Exclusive: Cisco plans hybrid NAC scheme

The company's forthcoming oneNAC technology is expected to address customer concerns about the complexity, maintenance and speed of Cisco's current security options. **Page 12.**

The metal whiskers threat

How concerned should you be about these data center intruders? **Page 24.**

NETWORKWORLD

The leader in network knowledge ■ www.networkworld.com September 3, 2007 ■ Volume 24, Number 34

Those crazy kids

The issue of dealing with Gen Y technology demands will be among those debated at next week's Security Standard event. **Page 14.**

Getting ready for multimegabit mobility

Columnist Johna Till Johnson on how you need to prepare. **Page 18.**

Cisco's Self-Defending Network strategy

Latest on this work in progress. **Page 26.**

Clear Choice Test: Corporate IM

IBM Lotus Sametime wins corporate instant message test. Jabber and Cisco are close seconds. **See full review, Page 42.**

Sharing the secrets of vendors' pricing plans

BY JON BRODKIN

Ever wonder why a software license costs as much as it does? If you suspect vendors charge as much as they can get, you wouldn't be far from the truth.

"It's primarily market-based," says Sally Bament, vice president of marketing at BlueNote Networks, which sells IP telephony software. It's inexpensive to manufacture software, so it doesn't make sense to base pricing on the vendor's cost, she says. R&D costs are taken into account in the company's overall business plan, so that doesn't play much of a role in pricing, either, she adds. If a software component is licensed from another vendor, the cost is taken into account in pricing but it is still a very small part of the equation.

See Pricing, page 20

MORE PRICING STORIES

- Why total cost of ownership is more important than you think. **Page 20.**
- **Vendor beware:** This CTO knows — and will exploit — your weaknesses. **Page 22.**

MPLS proposal spawns standards body turf war

IETF warns of possible Internet 'train wreck'

BY CAROLYN DUFFY MARSAN

The Internet's leading standards bodies are sparring over a set of next-generation network-transport specifications that some say could lead to massive interoperability issues for service providers if they are left unchanged.

The IETF is at odds with the International Telecommunication Union over a special transport network architecture the ITU's Telecommunications standards division (ITU-T) is developing to let MPLS traffic run over an Ethernet backbone. Among the network equipment vendors that have been contributing to the development are Alcatel-Lucent, Ericsson, Fujitsu and Tellabs.

The problem, according to the IETF, is that the ITU's Transport-MPLS (T-MPLS) specification will not work with the billions of dollars in routers and switches that carriers have installed in recent years based on the IETF's MPLS standards.

"The situation is catastrophic," says Stewart Bryant, IETF liaison to the ITU-T on MPLS issues and a technical leader at Cisco. "There's a fundamental opportunity

See MPLS, page 16

UTM firewalls: *READY FOR THE ENTERPRISE*

TESTS SHOW UNIFIED THREAT MANAGEMENT APPLIANCES AREN'T JUST FOR THE SMB MARKET ANYMORE.

ENTERPRISE IT has shied away from UTM firewalls because the appliances can cause performance problems, are tricky to deploy in large networks and don't always match the quality of best-of-breed tools. But the latest generation of UTM devices are worth a look because they reduce complexity, simplify management and improve flexibility. Not to mention that they promote long-term cost savings. **Page 35.**

GIACOMO MARCHESI

RIPPING & REP
IS SO OLD

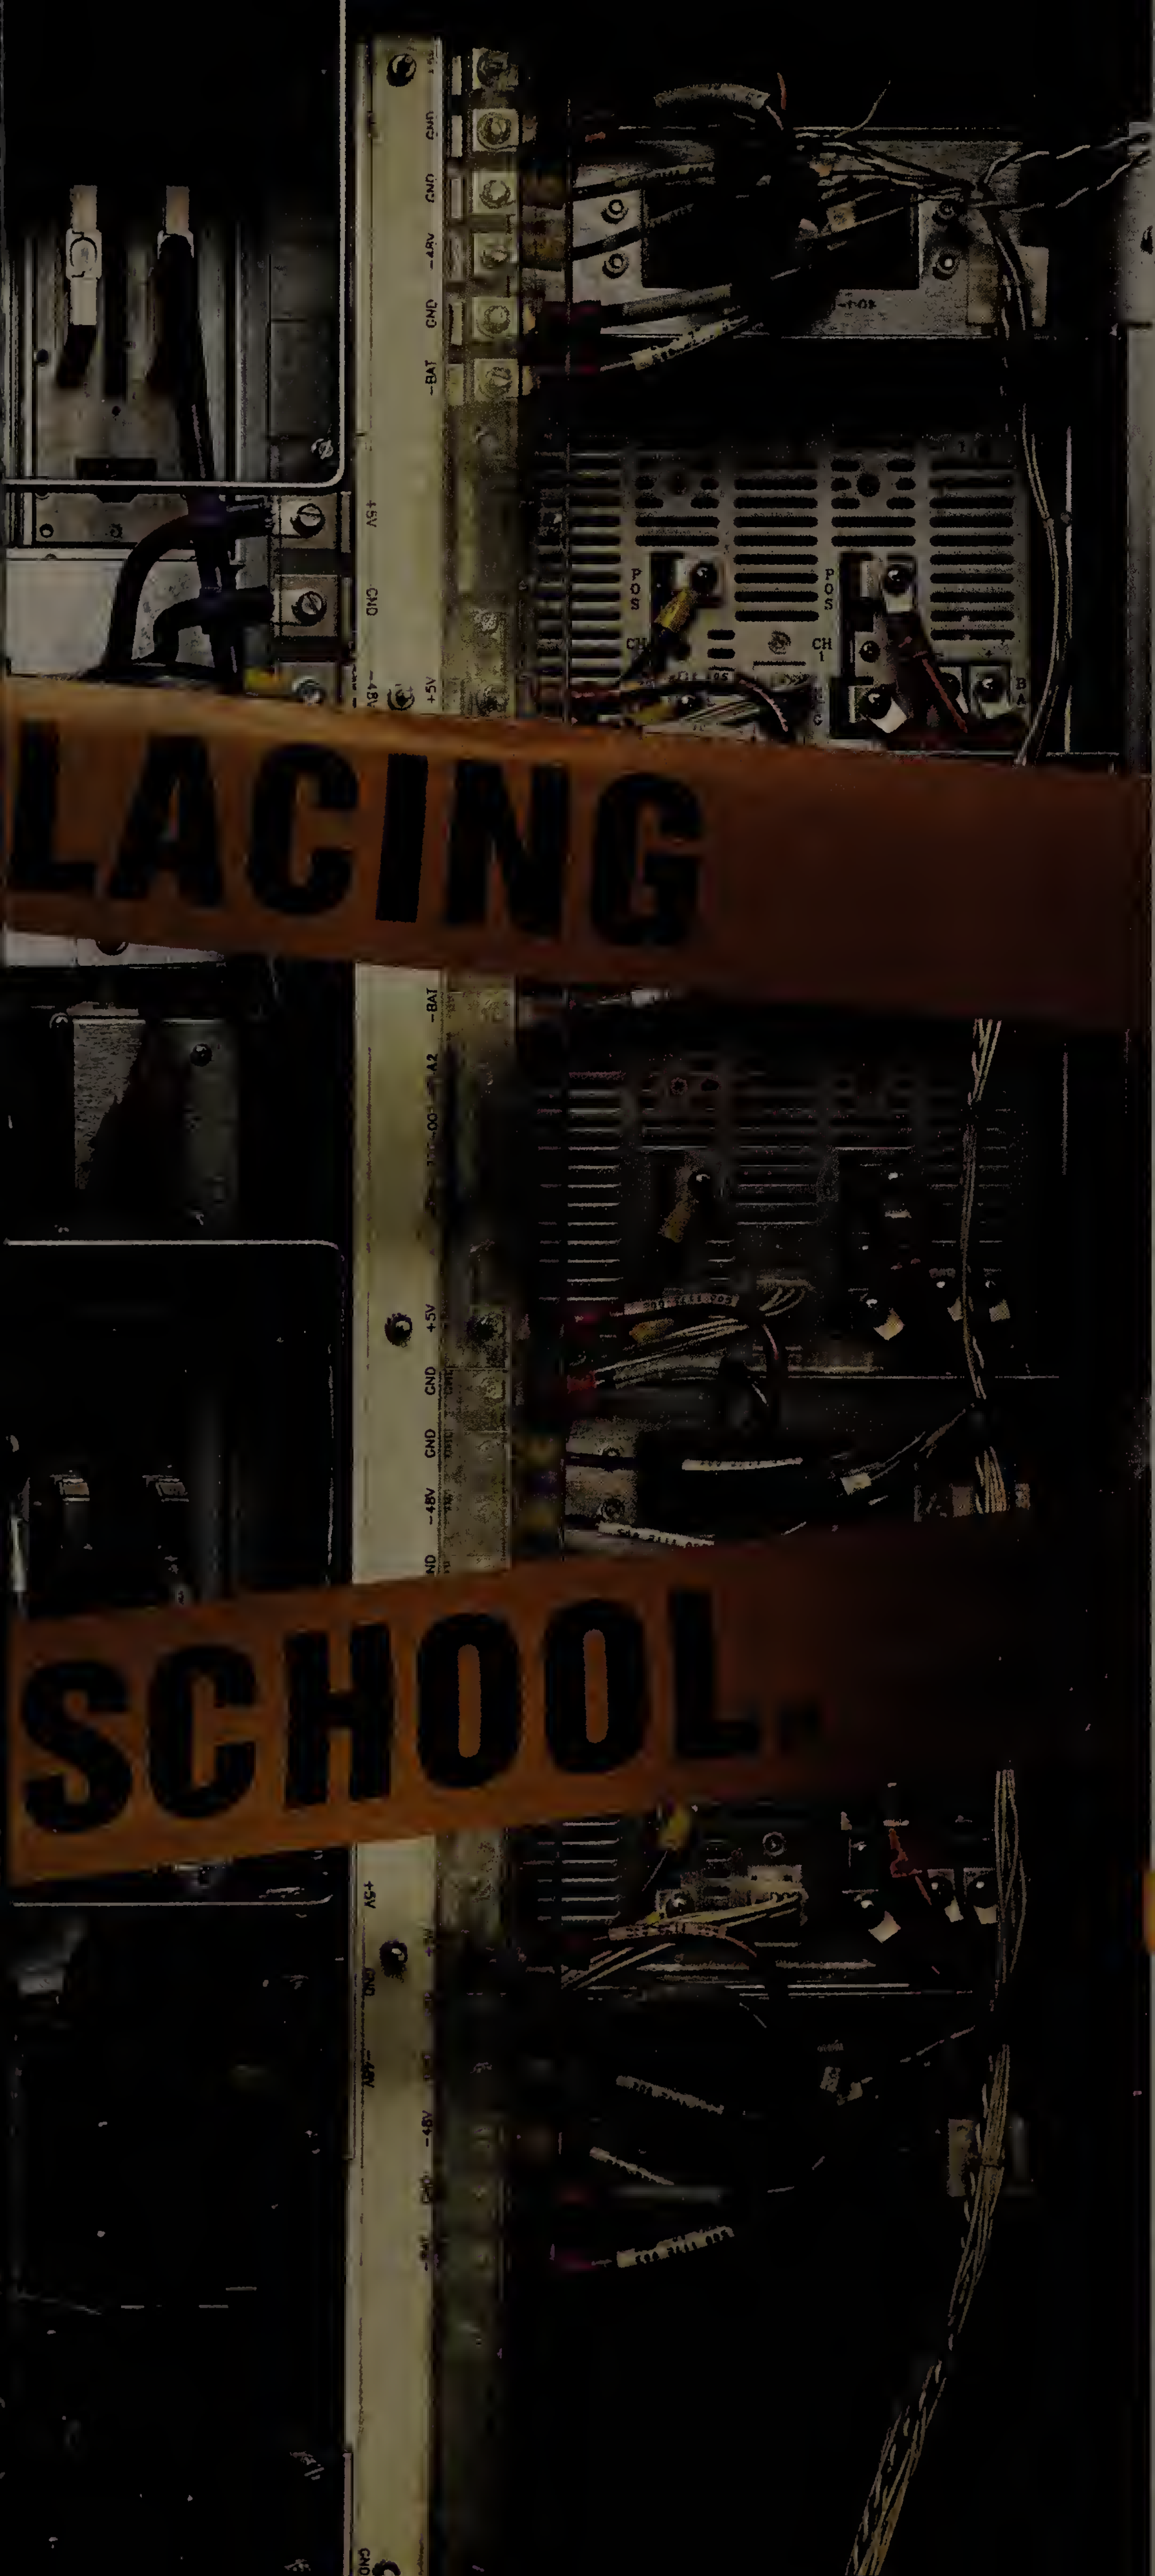

It's time to move to the future with the hardware you've got.

When moving to VoIP, ripping and replacing used to be the only way. Now, it's the out-of-date way. That's because it's no longer about hardware.

It's actually about software.

Now you can keep your hardware—your PBX, your gateways, even your phones. Simply move to VoIP with software. Software that integrates with Active Directory,® Microsoft® Office, Microsoft Exchange Server, and your PBX.

Maximize your current PBX investment and make it part of your new software-based VoIP solution from Microsoft. You're much closer to VoIP than you realize. Learn more at microsoft.com/voip

VoIP AS YOU ARE.

Your potential. Our passion.®

Microsoft®

09·03·07
NETWORK WORLD

UTM firewalls:

READY FOR THE ENTERPRISE

NETWORK INFRASTRUCTURE

ENTERPRISE COMPUTING

APPLICATION SERVICES

SERVICE PROVIDERS

TECH UPDATE

NETWORKWORLD.COM

■ **CONTACT** Network World, 118 Turnpike Road, Southborough, MA 01772; Phone: (508) 460-3333; E-mail: nwnews@nww.com; ■ **REPRINTS:** (717) 399-1900; ■ **SUBSCRIPTIONS:** Phone (508) 490-6444; E-mail: nwcirc@nww.com; URL: www.subscribenw.com

GOODBADUGLY

Goodbye, clutter

Massachusetts Institute of Technology researchers have come up with a way to measure visual clutter, a breakthrough that could help everyone from fighter pilots to Web site designers. The impetus for the work was that "we lack a clear understanding of what clutter is, what features, attributes and factors are relevant, why it presents a problem and how to identify it," says Ruth Rosenholtz, principal research scientist at MIT's Department of Brain and Cognitive Sciences.

China's cartoon cops

China Daily reports that two virtual police officers will start showing up on Beijing news portals and other Web sites to remind viewers about Internet security and give them an easy way to report illicit material. "They will be on the watch for websites that incite secession, promote superstition, gambling and fraud," the paper quotes an officer from the Beijing municipal public security bureau as saying.

More rootkits from Sony?

Sister publication *Computerworld* is reporting on a Finnish security company's claim that Sony is selling a line of USB drives that install files in a hidden folder accessible to and exploitable by hackers. The report brought back memories of Sony's ill-advised use of rootkit technology for copyright protection on CDs a couple of years back.

POLL

A snapshot of how networkworld.com visitors voted on a key networking issue last week:

What is your biggest worry about wireless data services?

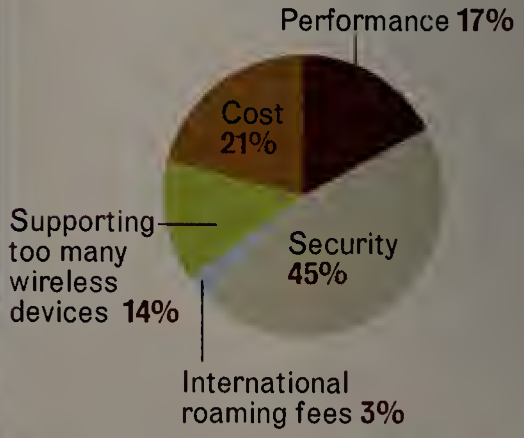

Vote and discuss: www.nwdocfinder.com/1460

Microsoft's antipiracy server meltdown

Re: WGA meltdown raises doubts about Microsoft reliability (www.nwdocfinder.com/1436):

I fully understand that there will be problems and issues with any system from time to time. However, when a malfunction causes PCs and users to cease to be productive —potentially rendering large numbers of machines to be unusable — that is just not acceptable.

For that matter, why is it necessary to constantly reverify that a copy of Windows is genuine? This should be something that could be engineered to happen once upon installation. If something happens that requires reinstallion, you validate again — one time. The requirement to validate your copy of Windows when an update or utility is downloaded from MS is patently ridiculous.

Ragtop

Discuss at www.nwdocfinder.com/1437

Risk management with wireless nets

Re: Just how wireless can we get? (www.nwdocfinder.com/1438):

Regarding jamming the wireless access point in a denial-of-service scenario, risk management should have taken into account that from the defense-in-depth viewpoint, an infrastructure that relies on connectivity through a single technology is asking for trouble. Just as router architectures and backbone designs take routing path redundancy and alternatives into account through whichever routing protocol is used, an access point should be doing the same. 802.11n is great if it's available; otherwise fall back to 802.11b/g and/or the the LAN/WAN connection. A WAP is not a huge device, whether it's plugged into a laptop /desktop PCMCIA port or part of a switch/router on the LAN.

The technology already exists in the form of the Cisco Mobile Access Router and other similar systems (via PC-104/PC-104+/PCI-104 technology) to create a Smart WAP or Smart WLAN/WWAN/WMAN architecture that can sense the loss of medium access (whether it's a backhoe that just severed a fiber-optic cable, or an radio frequency-based EMI event or DoS attack), address fall-back connectivity through the routing tables, reroute and press on. Whatever the event is that severs primary network connectivity, there should be a fall-back connectivity choice. Any network architecture that relies on a single technology for connectivity is a stationary target.

> “Any network architecture that relies on a single technology for connectivity is a stationary target.”

Bill Edwards

Discuss at www.nwdocfinder.com/1439

CEOs who game the system

Johnson writes about the CEOs who got greedy and got caught and now are in jail (Re: Another CEO falls . . . and few hear (www.nwdocfinder.com/1440). How about all the CEOs who, when they join a company, have a contract that no matter how they do, they win. I used to work for a worldwide data processing company. The CEO changed the severance package from two weeks for every year of service to two weeks for two years or less, and four weeks for more than two years. Then 45 days later there was a layoff of 1,500+. I knew of one worker that I had worked with who had 16.5 years of service and only got four weeks. But of course, when the CEO made a few mistakes — like the company losing $250 million when he tried to sell the company stock long and it went the wrong direction, a few contracts that he did not like that cost the company big-time, and then the stock went from high 60s to low teens. When he was let go, his package ran over $7 million per year — heads (good performance), he wins; tails (bad performance), the company and stockholders lose. I don't object when they get paid bonuses for doing good, but why make the rich when they perform badly?

You see this time and time again — I lose more faith in the system with this type of management than when someone goes to jail for violation of the law. When there is a violation of the law, it is usually pretty clear. This other type of management is a form of legal murder. But, today the Holy Grail of business is the bottom line and nothing else counts!

Joseph M. Brown

Discuss at www.nwdocfinder.com/1461

E-mail letters to jdix@nww.com or send them to John Dix, editor in chief, Network World, 118 Turnpike Road, Southborough, MA 01772. Please include phone number and address for verification

NETWORKWORLD

EDITORIAL DIRECTOR: JOHN GALLANT
EDITOR IN CHIEF: JOHN DIX

■ NEWS

ONLINE EXECUTIVE EDITOR, NEWS: BOB BROWN
ONLINE NEWS EDITOR: MICHAEL COONEY
ONLINE NEWS EDITOR: PAUL MCNAMARA
ONLINE ASSOCIATE NEWS EDITOR: ANN BEDNARZ (612) 926-0470

■ NET INFRASTRUCTURE

SENIOR EDITOR: JOHN COX (978) 834-0554
SENIOR EDITOR: TIM GREENE
SENIOR EDITOR: ELLEN MESSMER (941) 792-1061

■ ENTERPRISE COMPUTING

SENIOR EDITOR: JOHN FONTANA (303) 377-9057
SENIOR EDITOR: DENI CONNOR (512) 345-3850
SENIOR WRITER: ROBERT MULLINS (408) 306-4309

■ APPLICATION SERVICES

NATIONAL CORRESPONDENT: CAROLYN DUFFY MARSAN, (317) 566-0845
SENIOR EDITOR: DENISE DUBIE
SENIOR EDITOR: CARA GARRETSON (240) 246-0098
REPORTER: JON BRODKIN

■ SERVICE PROVIDERS

MANAGING EDITOR: JIM DUFFY (716) 655-0103

■ NET.WORKER

EDITOR: MICHAEL COONEY

■ COPY DESK/LAYOUT

MANAGING EDITOR: RYAN FRANCIS
COPY CHIEF: TAMMY O'KEEFE
SENIOR COPY EOITOR: JOHN DOOLEY
SENIOR COPY EOITOR: BOB SPRAGUE
COPY EOITOR: CAROL ZARROW

■ ART

DESIGN DIRECTOR: TOM NORTON
ART DIRECTOR: BRIAN GAIDRY
ASSOCIATE ART DIRECTOR: STEPHEN SAUER

■ FEATURES

FEATURES EOITOR: NEAL WEINBERG

■ CLEAR CHOICE TESTS

EXECUTIVE EOITOR, TESTING: CHRISTINE BURNS, (717) 243-3686
LAB ALLIANCE PARTNERS: JOEL SNYDER, Opus One; JOHN BASS, Centennial Networking Labs; BARRY NANCE, independent consultant; THOMAS POWELL, PINT; Miercom; THOMAS HENDERSON, ExtremeLabs; TRAVIS BERKLEY, University of Kansas; DAVID NEWMAN, Network Test; CHRISTINE PEREY, Perey Research & Consulting; JEFFREY FRITZ, University of California, San Francisco; JAMES GASKIN, Gaskin Computing Services, MANDY ANDRESS, ArcSec; RODNEY THAYER, Canola & Jones
CONTRIBUTING EDITORS: DANIEL BRIERE, MARK GIBBS, JAMES KOBIELUS, MARK MILLER

■ NETWORKWORLD.COM

EXECUTIVE EDITOR, ONLINE: ADAM GAFFIN
SITE EOITOR: JEFF CARUSO, (631) 584-5829
SITE PRODUCER: JONATHAN SUMMEY
SR. ONLINE EOITOR: LINDA LEUNG, (415) 341-1367
COMMUNITY EOITOR: JULIE BORT (970) 482-6454
PROGRAM DIRECTOR: KEITH SHAW, (508) 490-6527
MULTIMEDIA EDITOR: JASON MESERVE
ASSOCIATE ONLINE EDITOR: SHERYL HODGE
ONLINE ASSOCIATE ART OIRECTOR: ZACH SULLIVAN

■ NEW DATA CENTER SUPPLEMENTS

EDITOR: BETH SCHULTZ, (773) 283-0213
COPY EDITOR: CAROL ZARROW
EDITORIAL OPERATIONS MANAGER: CHERYL CRIVELLO
EDITORIAL OFFICE AOMINISTRATOR: PAT JOSEFEK
MAIN PHONE: (508) 460-3333
E-MAIL: first name_last name@nww.com

intel
Leap ahead™

NETWORKWORLD.COM

Follow these links to more resources online

BLOGOSPHERE

WSJ article advises workers on how to break IT policies. In herTech Exec blog, Linda Musthaler writes: "I just read the worst article ever on the Wall Street Journal Online edition. The July 30, 2007, edition of the Office Technology column tells non-IT workers how to get around the limits and policies that IT sets for office workers. . . . You've got to read the article to believe it, and when you do, you will be angry — very angry." **www.nwdocfinder.com/1448**

Your pipe is big enough. Cisco Subnet blogger Michael Morris asks: "Do you understand how much bandwidth 1Gbps is? That's the question I often find myself asking users. At our largest sites we have deployed Gigabit to the desktop — not because users need it, but because the price difference between 10/100/1000 and 10/100 cards in Cisco 6500s is small. . . . After several minutes of trying to explain to no avail to users they don't need 1Gbps since their network traffic patterns are not intense or sustained enough . . . I find my self asking them 'Do you understand how much bandwidth 1Gbps is? You don't need it. 100Mbps is fine!' "
www.nwdocfinder.com/1449

Microsoft Subnet welcomes two new bloggers. Chris Dalby is now writing the Essential Microsoft IT Toolkit blog. The blog covers products, cool third-party add-ons and his life as founder and director of Yellow Park, a Microsoft Certified Partner located in Kent, U.K. Dalby is known for his outspoken comments and his hilarious observations on life.
www.nwdocfinder.com/1450

David Platt is now writing the Why Software Sucks blog as the August guest blogger. Platt penned a book of the same name. He runs Rolling Thunder Computing, an education and consulting practice, and teaches software development at Harvard University Extension School. In 2002, Microsoft designated him a Software Legend. Microsoft Subnet has 15 free copies of his latest book to give away.
www.nwdocfinder.com/1451

Cisco Subnet welcomes new blogger. Wendell Odom joins as the August guest blogger. Odom is one of the most respected Cisco trainers around. He splits his time between writing books for Cisco Press and teaching classes for Skyline ATS. Cisco Subnet has 15 copies of his latest book to give away, too.
www.nwdocfinder.com/1452

ITVIDEO

INTERVIEWS, THE COOLEST TOOLS AND MORE

COOL TOOLS:

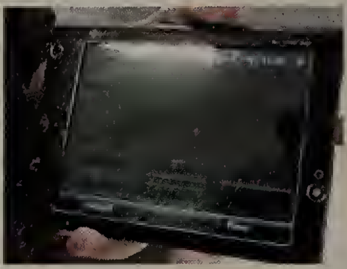

Not much Advantage here

The HTC Advantage 7501 promises a mobile office device with lots of bells and whistles. But can it survive in an iPhone world?

www.nwdocfinder.com/1445

TWISTED PAIR PODCAST:

Sleepless in the 'Twisted Lair'

Jason Meserve and Keith Shaw talk about whether a 17-year-old deserves a new car for hacking the iPhone, and explore the reasons why Keith isn't getting a lot of sleep these days.

www.nwdocfinder.com/1446

PANORAMA PODCAST:

SOX at 5: Benefits and headaches

James Sayles from Ecora Software talks with Cara Garretson about the five-year anniversary of the Sarbanes-Oxley Act.

www.nwdocfinder.com/1447

BEST OF NW'S NEWSLETTERS

Ready to abandon copper?

Microsoft has a ways to go with VoIP

Wide-area networking: It was recently reported in the popular press that in some cases, service providers are removing copper wiring capabilities when homes convert to fiber optic services. In this particular case, the situation involved a person ordering Verizon's FiOS fiber optic service. In this residential setting, the copper is being retired, and, according to an Associated Press article, future residents of the home may not have the option of going back to copper wiring. According to the article: "Under the Telecommunications Act of 1996, incumbent phone companies like Verizon must lease to rivals their copper network. That's generally not the case for next-generation fiber systems. And so far, Verizon has filed more than 100 notices with the FCC to retire portions of copper throughout its network."
www.nwdocfinder.com/1442

Convergence & VoIP: As we disclosed last time in our VoiceCon 2007 highlights, Microsoft has announced the addition of voice "quality of experience" monitoring to its unified communications and VoIP feature set. The timing of Microsoft's announcement was somewhat ironic because it came in the same week that Cisco CEO John Chambers and Microsoft CEO Steve Ballmer held a major press event in which they discussed how the two companies' relationship is formed around both cooperation and competition. Microsoft's announcement clearly reinforces the notion that it is a formidable competitor to Cisco and to other IP telephony equipment suppliers when it comes to VoIP and unified communications. Microsoft's Office Communications Server 2007 Quality of Experience Monitoring Server is designed to monitor voice and video quality and it features detailed analysis of network performance based on the user's endpoint.
www.nwdocfinder.com/1443

Network/systems management: Security information management products began to emerge earlier this decade as an alternative to manually dealing with the volume of security alerts generated across various network and security devices. A flurry of start-ups emerged — such as netForensics, GuardedNet, e-Security and Intellitactics to name just a few — with technology designed to marry the data collection, normalization and correlation capabilities of management software with the intelligence of security tools. **www.nwdocfinder.com/1444**

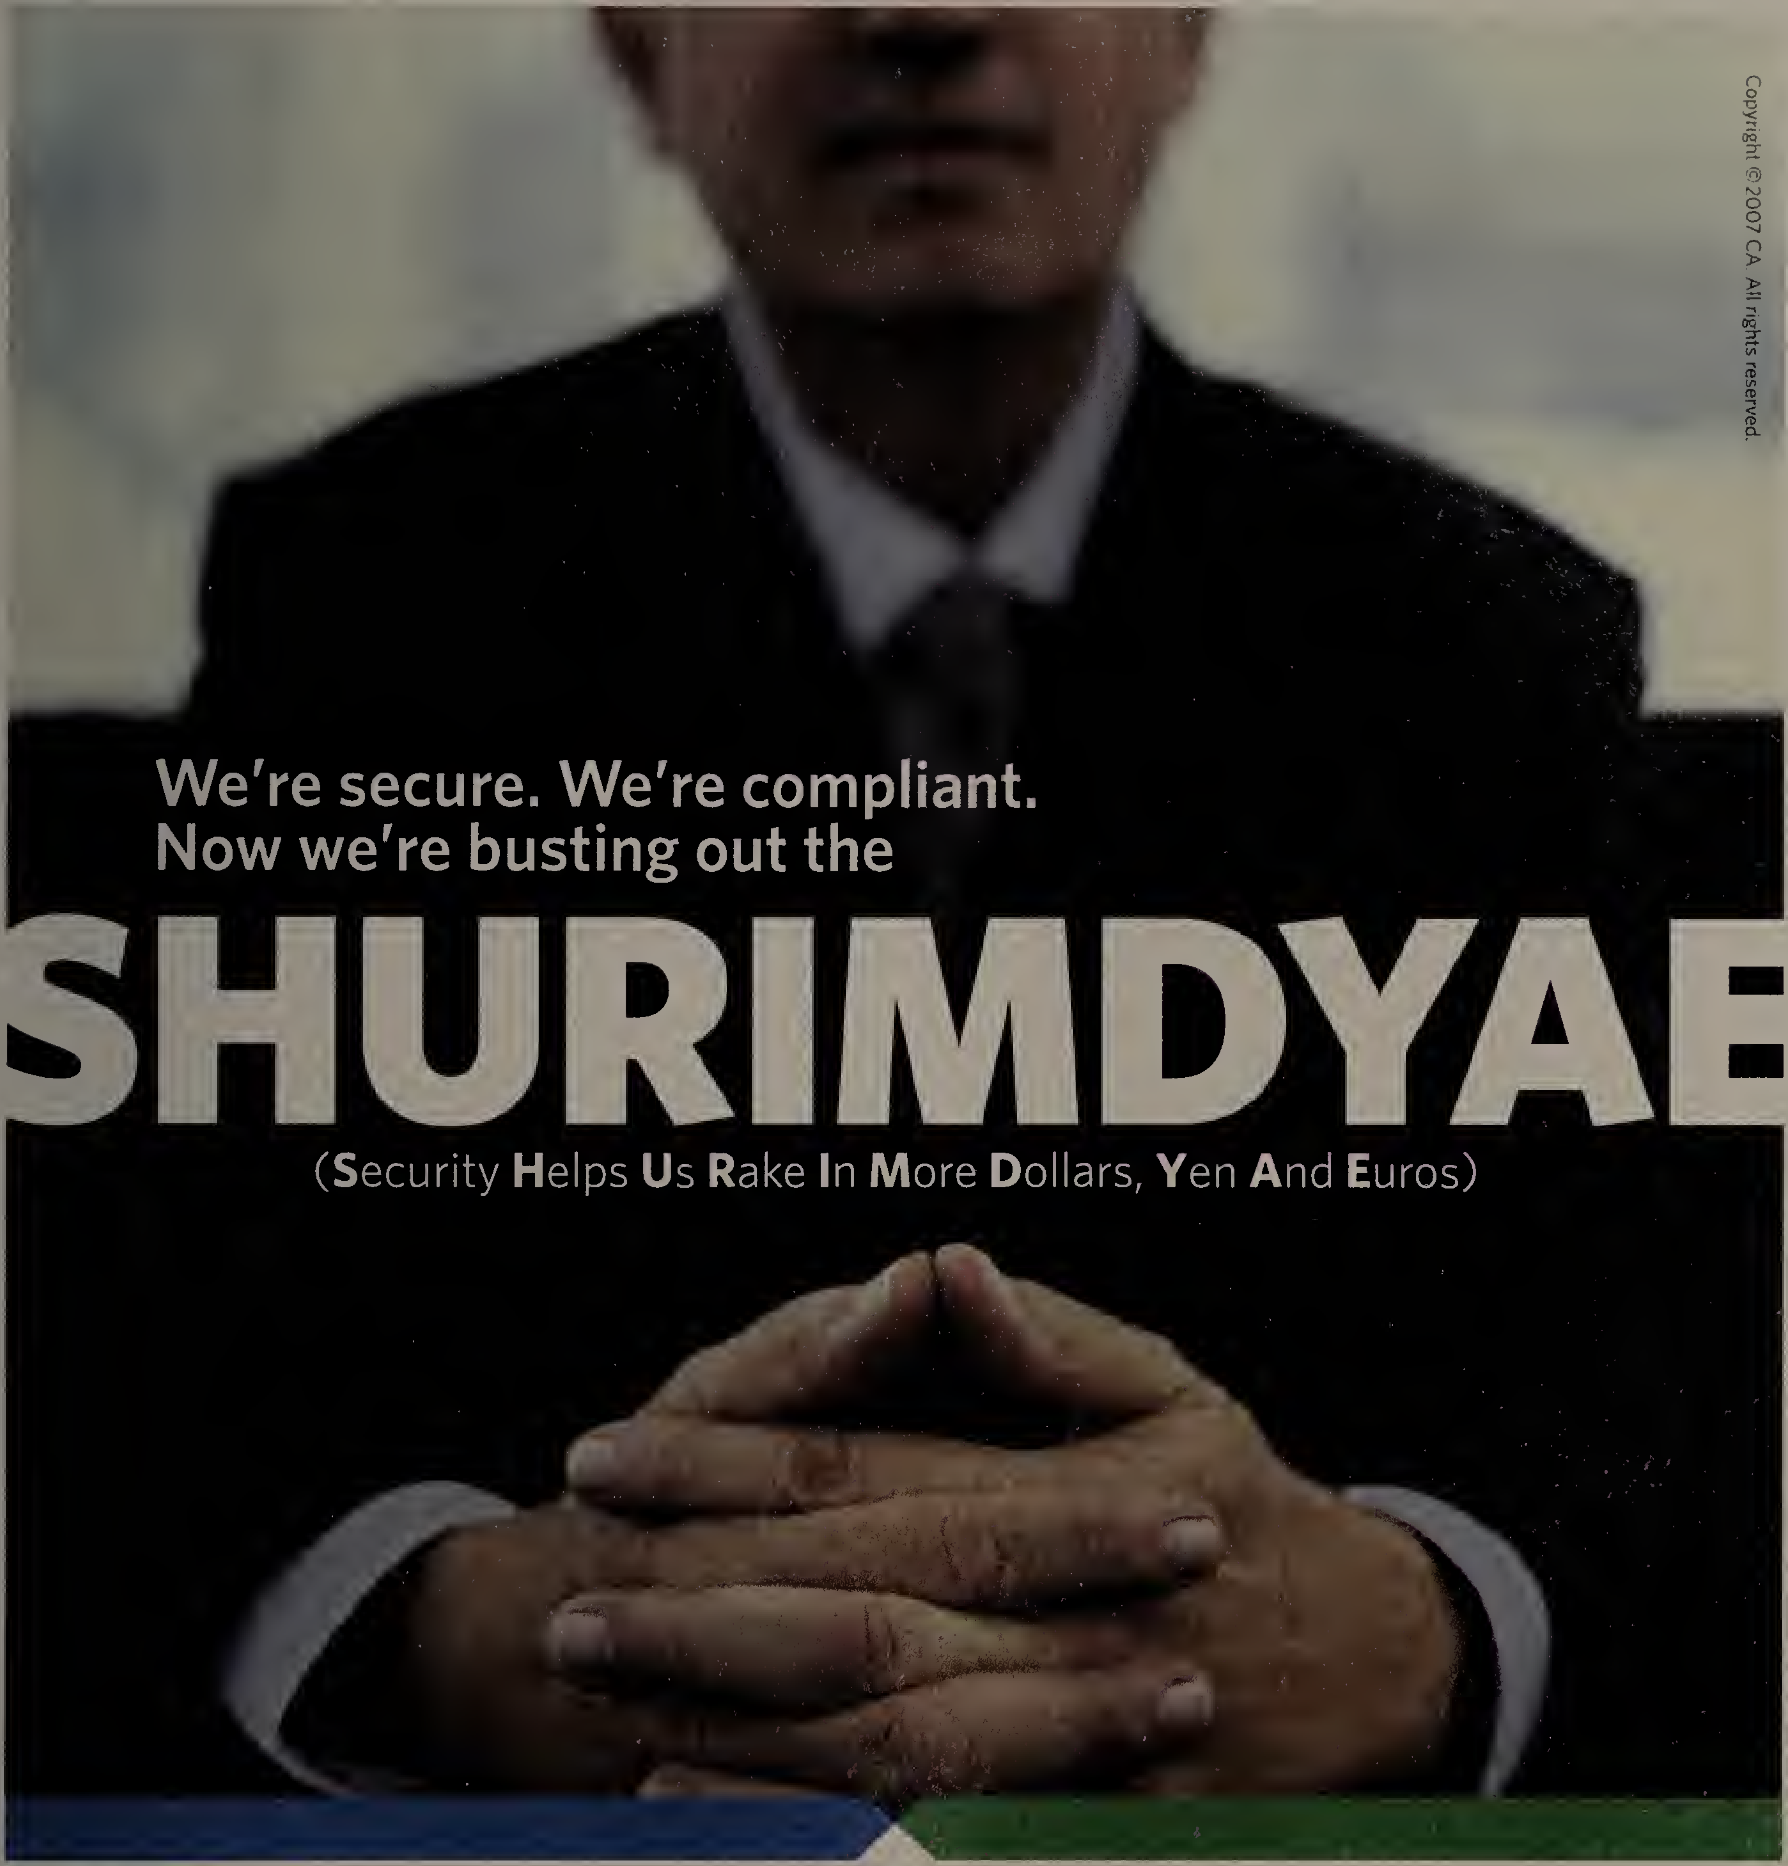
We're secure. We're compliant.
Now we're busting out the
SHURIMDYAE
(Security Helps Us Rake In More Dollars, Yen And Euros)

Microsoft acquires Parlano

Microsoft last week announced the acquisition of group-chat vendor Parlano, and said it plans to add the company's persistent-chat features to its portfolio of real-time communications wares. Persistent chat creates an ongoing instant messaging window that is organized with specific topics and can stretch across geographically dispersed workgroups. It includes security features, archiving capabilities and search tools.
www.nwdocfinder.com/1462

PDF spam levels plummet. It appears that PDF spam has had its 15 minutes of fame. Having reached its peak volume on Aug. 7 at nearly 30% of all spam messages sent, PDF spam today comprises less than 1% of spam, according to security vendor Sophos. One reason the unwanted e-mails with PDF files attached (usually pushing the recipient to purchase a penny stock) have all but disappeared is that e-mail users are starting to heed the warnings of IT managers that dictate attachments from unknown senders should not be opened.
www.nwdocfinder.com/1463

Sun powers start-up's Wi-Fi plans. A small U.S. start-up has announced technology for running Wi-Fi routers in remote places using only the power of the sun. Among the first round of products from Solis Energy is the Solar Power Plant, touted as being capable of supplying 12, 24 and 48 volts of DC power for use in such applications as surveillance cameras and outdoor Wi-Fi. Comprising a large solar panel connected to a generator, the system is said to be able to power such devices for as long as seven days without sunlight. The company also has a separate "tap adaptor" that can feed 120 volts of AC power to Wi-Fi, WiMAX and other outdoor systems from ordinary street lights.
www.nwdocfinder.com/1464

Microsoft sics lawyers on popular AutoPatcher utility. Microsoft last week shut down a popular utility built and maintained by Windows enthusiasts for installing updates offline. The AutoPatcher utility, created by project manager Antonis Kaladis, provides an interface to a large collection of updates, common applications and registry tweaks. The collection could be downloaded once, then used to update many computers, saving time and bandwidth. Microsoft, however, told Kaladis that it fears his utility potentially could distribute malicious software along with legitimate Microsoft updates.
www.nwdocfinder.com/1465

The benefits and shortcomings of NAC

On Aug. 28, security guru Joel Snyder conducted a live, online text chat on the topic, "The truth about NAC." What follows is a partial transcript of the event. The full transcript can be found at www.nwdocfinder.com/1458.

Network World chats

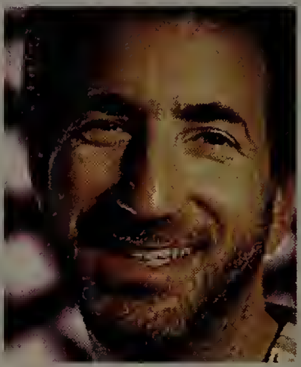

What's the biggest shortcoming you see with [network access control] implementations?

The lack of standardization of NAC approaches and strategies is really holding us back. We want to have different products for different requirements, but NAC products are so different across the board that it makes it difficult for people to know what will solve their needs. You have to be a product-evaluation guru just to understand some of the subtle differences between these products. I think that this will shake out over time, but if you look at [Mandy Andress' test of NAC alternatives] a few weeks ago, you'll see that she got really different products with really different designs (www.nwdocfinder.com/1466). This makes it hard to know what's right for you.

What are your thoughts about in-band versus out-of-band NAC solutions?

I'll have to throw a definition here, and see if you agree: in-band I think of as a box, like maybe a Vernier/ConSentry/Nevis or even Cisco CCA (in in-line mode, which is one option), which controls all access. Out-of-band is what I like to call "edge enforcement," more like 802.1X. Hybrid is more half-way, like Lockdown or CCA in that mode. Anyway, given those definitions: Edge is really where I think we want to go for big enterprise deployments. It scales, it handles the load and it doesn't depend on a single point to do enforcement. In-band I think of more for the occasional guest access — drop one of those boxes in between your guests and let it handle that load. Bam, problem solved, that was easy, etc. Of course, that doesn't mean that the in-band guys can't handle the load, but you really want to aim for edge enforcement if it fits, and go for in-band if it doesn't. And there are zillions of places where in-band fits better.

Should users hold off on implementing NAC until the vendors sort it all out?

Of course not. You need to buy, buy, buy, so those poor guys can keep up payments on their Boxsters. No, seriously, though, you can solve a lot of point problems with current solutions today and look to the future for better solutions with wider scope. I see a lot of people with pain points that need solutions — they should be going for something today. And a little experience today will help you pick the right solution tomorrow. Should you buy a NAC solution for 50,000 enterprise users on a Windows domain in 30 buildings? Well, I'd do a test rollout for a while first, if I were you.

What's your vision of NAC products five years from now?

Universal ho-hum. Just like VPN. We all have it where we need it and it's not so exciting. That's what we want. Universal dullness. We have to go to Funky Town, and then move to Dullsville. That's a good sign.

If NAC is ho-hum in five years, what in security is exciting in five years?

Dude. I'm going to be running a BBQ stand in five years. You call me up and tell me. ■

ONLINE: Enter the discussion

Upcoming chats feature Michael Osterman demystifying enterprise messaging and Amazon.com CTO Werner Vogels discussing the road to infinite capacity. See an archive of our first three chats.

www.nwdocfinder.com/1457

Cisco plans to blend its NAC schemes

OneNAC takes the best of its NAC Appliance and its network-based NAC

BY TIM GREENE

Cisco is planning a hybrid of its NAC architectures that will address customer concerns about the complexity, maintenance and speed of the company's current options.

The upgrades would make it possible for customers to buy Cisco's NAC Appliance — the NAC option most of its customers choose first — and later migrate to its network-based NAC Framework architecture without having to swap out as many elements.

NAC Appliance and NAC Framework now use different client software to evaluate the security posture of network endpoints as part of the NAC process. NAC Framework relies on its Access Control Server (ACS) to determine which access policy to apply, while NAC Appliance relies on its separate management server to determine if endpoints are in compliance.

Cisco calls the more unified NAC picture oneNAC, according to a source who knows what Cisco is saying to its customers about its NAC road map and who spoke on the condition of anonymity because the source's employer didn't authorize speaking to the press.

One of the problems all NAC customers face is that NAC appliances don't scale enough to accommodate a corporatewide deployment without using many appliances, says Rob Whiteley, an analyst with Forrester Research. The solution is network-based NAC, which scales for large deployments without requiring a proliferation of new devices on the network, he says. A migration strategy between appliance- and network-based NAC would simplify customers' transitions to wider NAC deployments.

Cisco describes its NAC plans as a path for customers to buy its NAC Appliance now and migrate to its NAC Framework over time.

"Our customers like to start with NAC Appliance because it's easier and doesn't require upgrading their infrastructure gear all at once, but they also like many aspects of the Framework approach," a Cisco spokesman said in an e-mail. "So, in interpreting 'oneNAC', it refers to making sure both solutions are interoperable with each other, that customers get investment protection, etc. That way, customers can upgrade infrastructure as part of the natural refresh cycle while getting started with NAC."

Cisco's NAC Appliance sits inline with traffic to enforce access policies. The throughput is 1Gbps, a limiting factor for faster networks. The appliance also can be deployed out of the traffic stream — out-of-band — and use Cisco network switches to enforce NAC policies.

Cisco Framework relies on software deployed on network endpoints in combination with Cisco's ACS/RADIUS server to trigger 802.1X enforcement of admission policies. One drawback customers find is that adding and updating policies is complex because it involves directly touching the RADIUS server and refreshing local policy directories, the source says. "The technology is there, but to get the implementation is a battle," the source says. OneNAC would draw on pieces of both architectures. It would use the management-server portion of the NAC Appliance implementation as the single place for customers to create, add and change NAC policies, and it would be fully compliant with the 802.1X authentication standard, the source says.

The new flavor of Cisco NAC also would consolidate NAC client software that reports on the configuration of endpoints. The Cisco appliance- and Cisco's network-based NAC products use different clients, and oneNAC would create a single client that serves both scenarios, the source says.

Cisco's oneNAC is 12 months to 18 months from being available, the source says.

Cisco has an advantage in that it owns its own RADIUS server technology and can customize its interactions freely with its NAC platform. Among its competitors, only Juniper Networks, with its Steel Belted Radius server, owns its own RADIUS technology.

It is very possible to deploy NAC that relies on standard interfaces with RADIUS servers, as has been demonstrated at Interop.

Unlike smaller vendors that sell appliances that work within existing networks, Cisco makes the switches that are used as enforcement points, making customization and extended features a possibility. ■

Cisco's NAC migration

Cisco plans to blend parts of its two network-access-control architectures so customers can buy one and gradually shift to the other while still using the initial purchase.

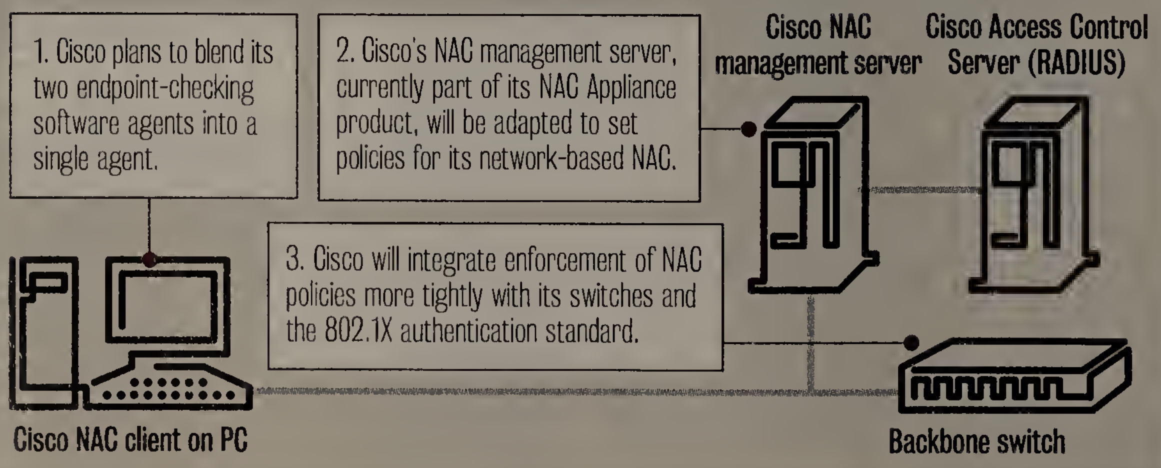

1. Cisco plans to blend its two endpoint-checking software agents into a single agent.

2. Cisco's NAC management server, currently part of its NAC Appliance product, will be adapted to set policies for its network-based NAC.

3. Cisco will integrate enforcement of NAC policies more tightly with its switches and the 802.1X authentication standard.

Cisco NAC management server
Cisco Access Control Server (RADIUS)
Cisco NAC client on PC
Backbone switch

InBrief

Windows Server 2008 delayed again

Microsoft again delayed the release of Windows Server 2008, saying development of the software could take as long as another three months beyond its previously planned December release. The new plan calls for the server to be released to manufacturing between Jan. 1 and March 31, 2008. The delay is being blamed on the need for more testing of the server, which was first put into beta in 2005 and has suffered numerous delays and feature dumps. Windows Server 2008 focuses on three areas: management, including Server Core; security, such as BitLocker drive encryption and read-only domain controllers; and performance, including a redesigned TCP/IP stack.

Data breaches hurt corporate image; some customers loyal

Data breaches have a strong emotional impact on consumers but don't always lead them to abandon the company as a customer, according to a survey sponsored by data security vendor Tablus. Although 21% in the survey responded they had indeed stopped shopping at stores where confidential records had been stolen, 43% indicated they wouldn't. The remaining 36% didn't have a firm opinion about shopping at stores with a history of losing sensitive customer information.

"as a security measure, pogo sticks worked for awhile, then we opted for the stilts which worked brilliantly... until today."

Facing Generation Y security issues

Young employees entering the workforce bring a new round of security threats

This is the final story in a five-part series about the key security issues that will be addressed at The Security Standard event scheduled for Sept. 10-11 in Chicago.

BY CARA GARRETSON

As young adults who grew up on e-mail and online chat enter the workforce, they bring with them a set of newer technologies designed for rapid-fire communication and workplace personalization. Much of this technology may represent better, faster ways of getting a job done, but it also introduces a new round of security threats for corporate networks; and the decision to allow them or not must be made carefully.

These technologies — personal gadgets such as MP3 players, thumb drives, cell phones and PDAs; real-time communication technologies such as instant messaging and text messaging; and social-networking Web sites such as Facebook and MySpace — are part and parcel of the young workforce today, experts say. Called Millennials or Generation Y, this group is defined loosely as having been born between 1977 and 2002, and totals 70 million Americans — a large percentage of whom are bound to have one of the 100 million iPods sold to date in their pocket.

Many Generation Y technologies may offer an improvement over today's status quo — an IM or text message is likely to get the recipient's attention more quickly than an e-mail that sits in an in-box — but they can introduce serious security threats to corporate networks, according to some security vendors (see graphic).

For example, "the newer forms of attacks take advantage of Web sites with rich content and features: AJAX-enabled applications, embedded JavaScript and so on. These aren't really new technologies, but they're more pervasive now," says Paul Ferguson, network architect at Trend Micro. "And with [components like] Google Maps, where the processing is done on the PC instead of on the Web page, criminals are exploiting that avenue of content delivery. The ability for Web 2.0 applications to deliver that content is a Catch-22, because it also can allow you to be exploited."

For security professionals, it may seem that the prudent thing to do is to disallow the use of this kind of technology in the workplace: blacklist non-business-related Web sites; ban handheld or pocket devices; require employees to use company-issued and maintained laptops, PDAs and cell phones. After all, as much as 40% of employee Internet activity is non-work-related, according to IDC.

Experts warn, however, that such stringent policies can have a negative effect on the workforce and its productivity, as well as the company's ability to attract and keep valued workers. "It's part of the way [young employees] have grown up, part of what they expect," says Tony Kerns, deputy managing partner with Deloitte & Touche. "The global pressure on the workforce right now is huge; people are drawn all over the world by great, interesting offers that are not just money but also a lifestyle."

Earlier this year, security vendor MessageGate, which makes e-mail management software and was spun out of Boeing in 2003, conducted a series of roundtable discussions with senior IT professionals and young adults entering the workforce to try to understand the issues around Generation Y technology.

One thing MessageGate learned is that younger workers' preferences for newer technology often can be good news for an organization's IT department, according to Robert Pease, the company's vice president of marketing.

"When [older workers] first entered the workforce, we could communicate with each other via e-mail, and there was a big blurring between business and personal," Pease says. Today, young workers would rather communicate with each other via text messaging or postings on Web sites, and are less inclined to misuse the corporate e-mail system with personal messages, he says. "There's a bit more discipline around corporate communications today. The bad news is, how do I control" the other channels of communication?

One risk manager at a large financial-services company who asked not to be named sees the value in providing employees with a flexible work environment, but says that flexibility must be accompanied by well-defined policies (see www.nwdocfinder.com/1429) and layers of security technology. "Whenever employees are given flexibility for their hours and environment, you'll definitely have a happier, as well as more productive workforce," the risk manager says. He adds, however, "you need to specifically define parameters for what is and is not allowed in your policies, and spell out what will be the result of any violations."

Companies that believe they have communicated their policies sufficiently might need to think again. According to a survey by security vendor Senforce last March, 73% of the 308 respondents said they store corporate data on removable media, and 46% said they did not have — or were unaware of — corporate security policies that protect that information.

Although presenting a flexible work environment would be particularly important for companies whose employees are their assets — advertising and design firms, for example — the need to maintain a happy workforce is important in any industry. "It needs to be presented as a win-win situation," the risk manager says. "Explain to the employees that following the guidelines will help to ensure the continued flexibility of the work environment. If you make things too restrictive, younger employees may just pack up and go elsewhere." ■

E-mail is so two hours ago

Here are some of the technologies often used by young employees that can cause problems in business settings:

Technology	Potential workplace problems
USB storage devices	Can be used to steal corporate data; enough capacity to take large amounts of information, but small enough to go undetected.
iPods and other MP3 players	Can be set to steal corporate data. Also, downloading music and video can clog bandwidth.
Instant messaging	Many public networks don't offer security features; often chats aren't logged so there is no audit trail or proof of the communication; the real-time nature of chat can disrupt the workplace.
Cell phone text messaging	No ability to send file attachments; no communication log or audit trail.
Web 2.0 sites	Popular social-networking and related sites rely on technology with weak security that hackers are targeting as agents for downloading malware.

MPLS

continued from page 1

for a major train wreck" between the IETF's MPLS and the ITU-T's T-MPLS.

Bryant says the problem is that T-MPLS uses the same EtherType as MPLS, which will lead to confusion in operational networks. An EtherType is a field in the Ethernet network standard that indicates which protocol is being transported. "If you think about a piece of network equipment, it looks at the EtherType and that tells it how to process the packet. The EtherType is the same for MPLS and T-MPLS, so we are extremely unhappy about that. T-MPLS should use a different EtherType, ideally called T-MPLS so there is absolutely zero confusion in dealing with T-MPLS or MPLS traffic," he says.

"Our concern is that there should be absolutely nothing designed, implemented or specified that risks the deployed base of MPLS equipment," Bryant adds.

ITU not worried

T-MPLS is being developed by the ITU-T's Study Group 15 Working Party 3, which considers optical transport network structure. This group has been developing T-MPLS for three years and has finished four specifications, including an architecture document, a network-to-network interface, an equipment specification and a switching document. The ITU-T says service providers need a special profile or subset of MPLS to meet their requirements and that's why they are developing T-MPLS. For example, T-MPLS will support more robust protection-switching and operating environments and messaging than is provided by the IETF's MPLS standards.

ITU-T leaders deny that T-MPLS will create interoperability problems for the Internet, let alone catastrophic ones.

"I have a pretty good degree of confidence that we haven't put anything into the T-MPLS standard so far that's going to cause massive interoperability problems," says Stephen Trowbridge, chairman of the working group. Trowbridge, who works for Alcatel-Lucent, says, "T-MPLS will stay in the service provider network, and the customer network doesn't use it."

Trowbridge calls the T-MPLS flap a "turf war" between the IETF and the ITU-T, and he says emotions are running high among the members of the two standards bodies.

"This is sort of a contentious area," Trowbridge says. "Everything is converging. You see more and more optical technology further toward the edge of the network. You see more and more packet technology moving toward service provider networks. . . . Turf battles are inevitable."

The T-MPLS working group will hold a week-long meeting in Stuttgart, Germany, beginning Sept. 10, which several IETF leaders will attend in an attempt to hammer out a solution. Around 40 representatives from carriers and network equipment vendors plan to attend. "We'll put the proposals on the table and have an open discussion and try to resolve the issues," he says. "A stalemate isn't good for anybody."

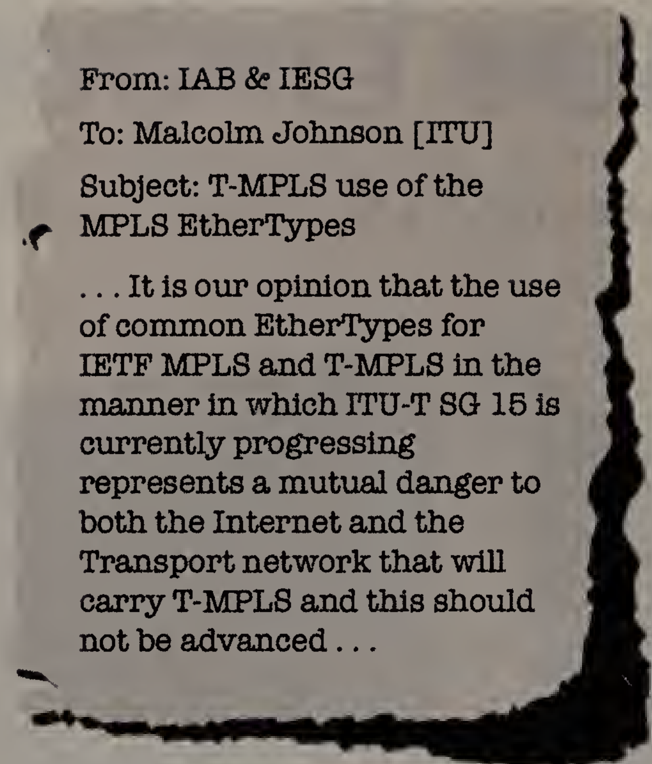

From: IAB & IESG

To: Malcolm Johnson [ITU]

Subject: T-MPLS use of the MPLS EtherTypes

. . . It is our opinion that the use of common EtherTypes for IETF MPLS and T-MPLS in the manner in which ITU-T SG 15 is currently progressing represents a mutual danger to both the Internet and the Transport network that will carry T-MPLS and this should not be advanced . . .

To read the complete letter, go to: **www.nwdocfinder.com/1459**

After the meeting is over, the ITU-T plans to send a letter to IETF leadership outlining the decisions the working party has reached regarding the direction of T-MPLS.

The ITU-T is using several IETF-developed technologies in T-MPLS, including MPLS EtherTypes and Pseudowire Emulation Edge to Edge for its codepoints. T-MPLS also duplicates the control, management and forwarding planes used by the IETF's MPLS standards.

The IETF charges that T-MPLS uses these technologies in a different and incompatible way from how they are defined in the IETF's MPLS standards. Therefore, T-MPLS and MPLS traffic cannot coexist on a network, the IETF says.

The ITU-T says this situation is fine because T-MPLS will be used only on service provider networks, not on enterprise networks.

"In order to carry an enterprise network's MPLS traffic over a service provider's T-MPLS network, the enterprise MPLS will go over Ethernet. It will be T-MPLS in the service provider network, and when it gets delivered to the customer, the T-MPLS label will be removed and the customer gets back traffic on top of Ethernet," Trowbridge explains. "When you look at that service model, it's hard to see how there could possibly be any protocol conflict."

The IETF's leadership considers that view unrealistic. "It is our experience that even with careful planning and design, network elements rarely remain disjoint in practice," the IETF leadership said in a strongly worded letter (www.nwdocfinder.com/1459) sent to Malcolm Johnson, director of the ITU's Telecommunication Standardization Bureau, in late July urging the Geneva-based standards body to change its course on T-MPLS. "Accidental configuration does occur and can be a significant factor in serious network outages and other problematic events."

The IETF is proposing the IETF and ITU work together to bring T-MPLS requirements into the IETF standards process to make sure they will work with the IETF's existing MPLS standards.

Alternatively, the IETF recommends that the ITU change T-MPLS so that it uses different codepoints in the control, management and forwarding planes.

This isn't the first time the IETF has raised the issue of how T-MPLS uses the same EtherTypes as MPLS. The IETF sent a letter to the ITU-T a year ago asking that T-MPLS use different EtherTypes. The IETF sent representatives to a meeting in France in September 2006, and the consensus of that meeting was to use common EtherTypes in MPLS and T-MPLS.

"It was our understanding that the EtherTypes issue was resolved," Trowbridge says, adding that his working group was surprised by the recent letter and rhetoric from the IETF leadership. "This letter of July 2007 was the first indication we had on the ITU side since our reply of September 2006 that they didn't consider the issue closed."

Thought problem was resolved

If the T-MPLS issue goes unresolved, service providers and enterprises rolling out MPLS technology may be harmed, Bryant says. "T-MPLS is designed to be deployed inside service provider networks," he says. "In as much as an enterprise uses a service provider for its infrastructure, then they clearly need to be concerned that those service provider networks are correctly functioning and providing the MPLS service they are looking for."

Some large enterprises such as government agencies could roll out T-MPLS directly. "T-MPLS could find itself deployed in an enterprise network that has its own transport," Bryant adds. Instead of using T-MPLS, carriers could run MPLS directly over Ethernet using the IETF's Pseudowires technology, he says.

Bryant says T-MPLS has caused more friction between the IETF and ITU than is normally involved in Internet standards development.

"This is unusual," Bryant admits. "There have been many discussions between the IETF and the ITU where we've tried to work together on this. We do hope that we can work together and that we can resolve this in an amicable way. We want to produce technology that satisfies everyone, but does it in such a way that there is no confusion going on in the network." ■

NEC IP Telephony Solutions

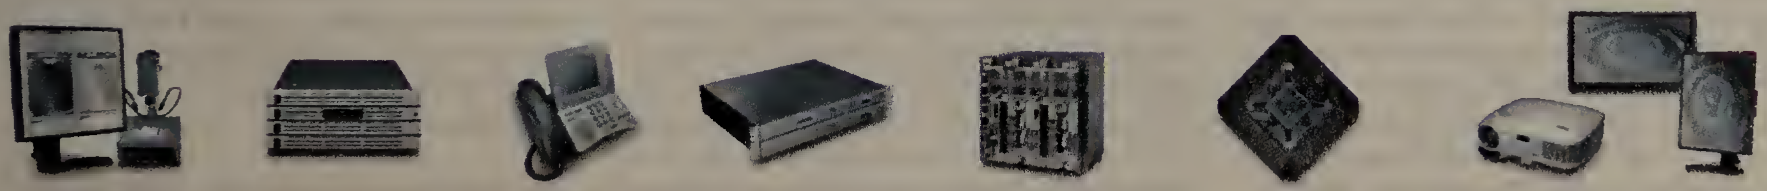

Get ready for multimegabit mobility

EYE ON THE CARRIER
Johna Till Johnson

Wireless data technologies have been coming of age for at least the past couple of decades. Remember CDPD? And it wasn't so long ago that Wi-Fi was new and exciting. Wireless data technologies seem to periodically "arrive" every decade or so — without ever managing to have a truly significant impact on more than a core group of users.

But all signs indicate that high-speed wireless data services are finally really arriving — and in a big way this time. For one thing, the sheer number of mobile users is skyrocketing. According to several research organizations, roughly 1 million new mobile subscribers come online in India every month. Many (perhaps most) of those are consumers, but enterprises that I've spoken with project 100% to 500% growth in the number of mobile-enabled employees (in all geographies) by mid-2008.

An increase in the number of mobile users is just part of the story. Even more significant is the increase in mobile bandwidth to each of those users. Mobile and wireless services are rapidly transforming from "poor man's connectivity" with data rates well below those for fixed services to comparable in speed and quality to their fixed-line counterparts.

By some projections, mobile broadband services will overtake fixed broadband services as early as 2010. And technologies such as HSPA and LTE deliver 1M to 10Mbps throughput to mobile users. That's enough to handle today's traffic mixes (e-mail, Web browsing, file transfer) as well as tomorrow's (interactive video, streaming multimedia).

What are the implications? For starters, IT executives need to stop thinking of wireless and mobile technologies as a niche — relevant for a subset of users but a footnote in the organization's overall strategy. Instead, they should assume that mobile connectivity will become an increasingly important piece of the technology road map, and plan and budget for it accordingly. That means rethinking current approaches to security and management, as well as revisiting overall costs (mobility adds significantly to per-employee IT costs). It also means envisioning ways in which business processes can be enhanced and improved.

ONLINE: Wireless LANs and enterprise mobility

Always-available access to information — and the ability to act on it instantly, anywhere — is an advantage in today's hyper-competitive world. Hone your edge. Join us on Sept. 6 at IT Roadmap: Dallas. Qualify to attend free at:

www.nwdocfinder.com/9159

More broadly, as I mentioned in last week's column, planners and legislators need to revisit global telecom policy in the context of emerging broadband wireless. Today, large chunks of spectrum are allocated to services such as analog TV that are virtually obsolete. And wireless technologies (including but not limited to GPS) can also potentially play a significant role in revised and enhanced emergency services.

Finally, network architects at enterprises and service providers need to rethink network designs as last-mile connectivity grows and evolves. Today, the typical user consumes a megabit per second or less in WAN connectivity. But as broadband wireless becomes the norm rather than the exception, applications will evolve to expect and consume much more, increasing performance requirements on edge, access, and core routers and switches.

The bottom line: Wireless has been around for so long we've begun to take it for granted — and that's a mistake. It's time to plan for tomorrow's multimegabit mobile networks.

Johnson is president and senior founding partner at Nemertes Research, an independent technology research firm. She can be reached at johna@nemertes.com.

Security-oriented architectures?

RISK & REWARD
Andreas Antonopoulos

SOA is one of those buzzword acronyms that mean so many things to so many people it's hard to pin down what it is. Nevertheless, many large enterprises are integrating applications and building applications using XML, Web services and rudimentary service-oriented architectures. But what about security?

An SOA is meant to provide enterprises with the means to develop applications rapidly by mixing together small, self-contained application services. What used to be "internal" communication in an application becomes an external network transaction. Because large enterprises are using these technologies and architectures already, we sought to learn to what degree enterprises have begun thinking about securing their SOA-based applications. The answer — very little. Just one-third are planning to implement SOA security within the next year.

Why the relatively low level of interest in SOA security? Quite frankly, companies still are getting their arms around how SOA-based applications will affect their overall architectures, not just security. SOA security is an issue on the horizon, but it's one of several.

"I'm worried about bots and botnets," says the head of security for a large university. "It seems to me that we're on the cusp of a new generation of attack tools that are precisely going to find vulnerability in these applications, much more so than they do now. Apps don't do a good job separating application from presentation layer. I'm imagining a scenario where agents look for and exploit very subtle vulnerabilities."

That said, SOA security is one area where companies at least are planning to put their money where their mouths are: 50% say they expect their SOA security budgets to increase during the next 12 to 18 months. That's not too difficult, given the low levels most folks are starting from: $78,000 was the mean spending of the handful of companies reporting they had an SOA security budget. Of course, there's also the question of what products companies are going to spend their money on. Leading-edge enterprises complain there's a lack of standardized products: "The mechanisms to date have not resulted in products that people are using. We have an initiative to look at message-brokering facilities. We have deployed XML gateways for security purposes. With [the Web Services Security protocol] we are not seeing much [vendor standards] agreement in that space," says an IT executive at a financial-services firm.

And, unsurprisingly, just a quarter of IT executives say they're using SOA-enabled devices in their security infrastructures.

The take-away? Mixed, but intriguingly so. Unlike the case with other communications-security issues (in particular, mobility and VoIP), IT executives seem to have aligned their SOA-security investment strategies with their priorities. As SOA activities in the enterprise continue to increase, we expect security budgets to follow. As I embark on further research in enterprise applications, I surely will be returning to this topic!

Antonopoulos is senior vice president and founding partner at Nemertes Research, a technology research firm. He can be reached at andreas@nemertes.com.

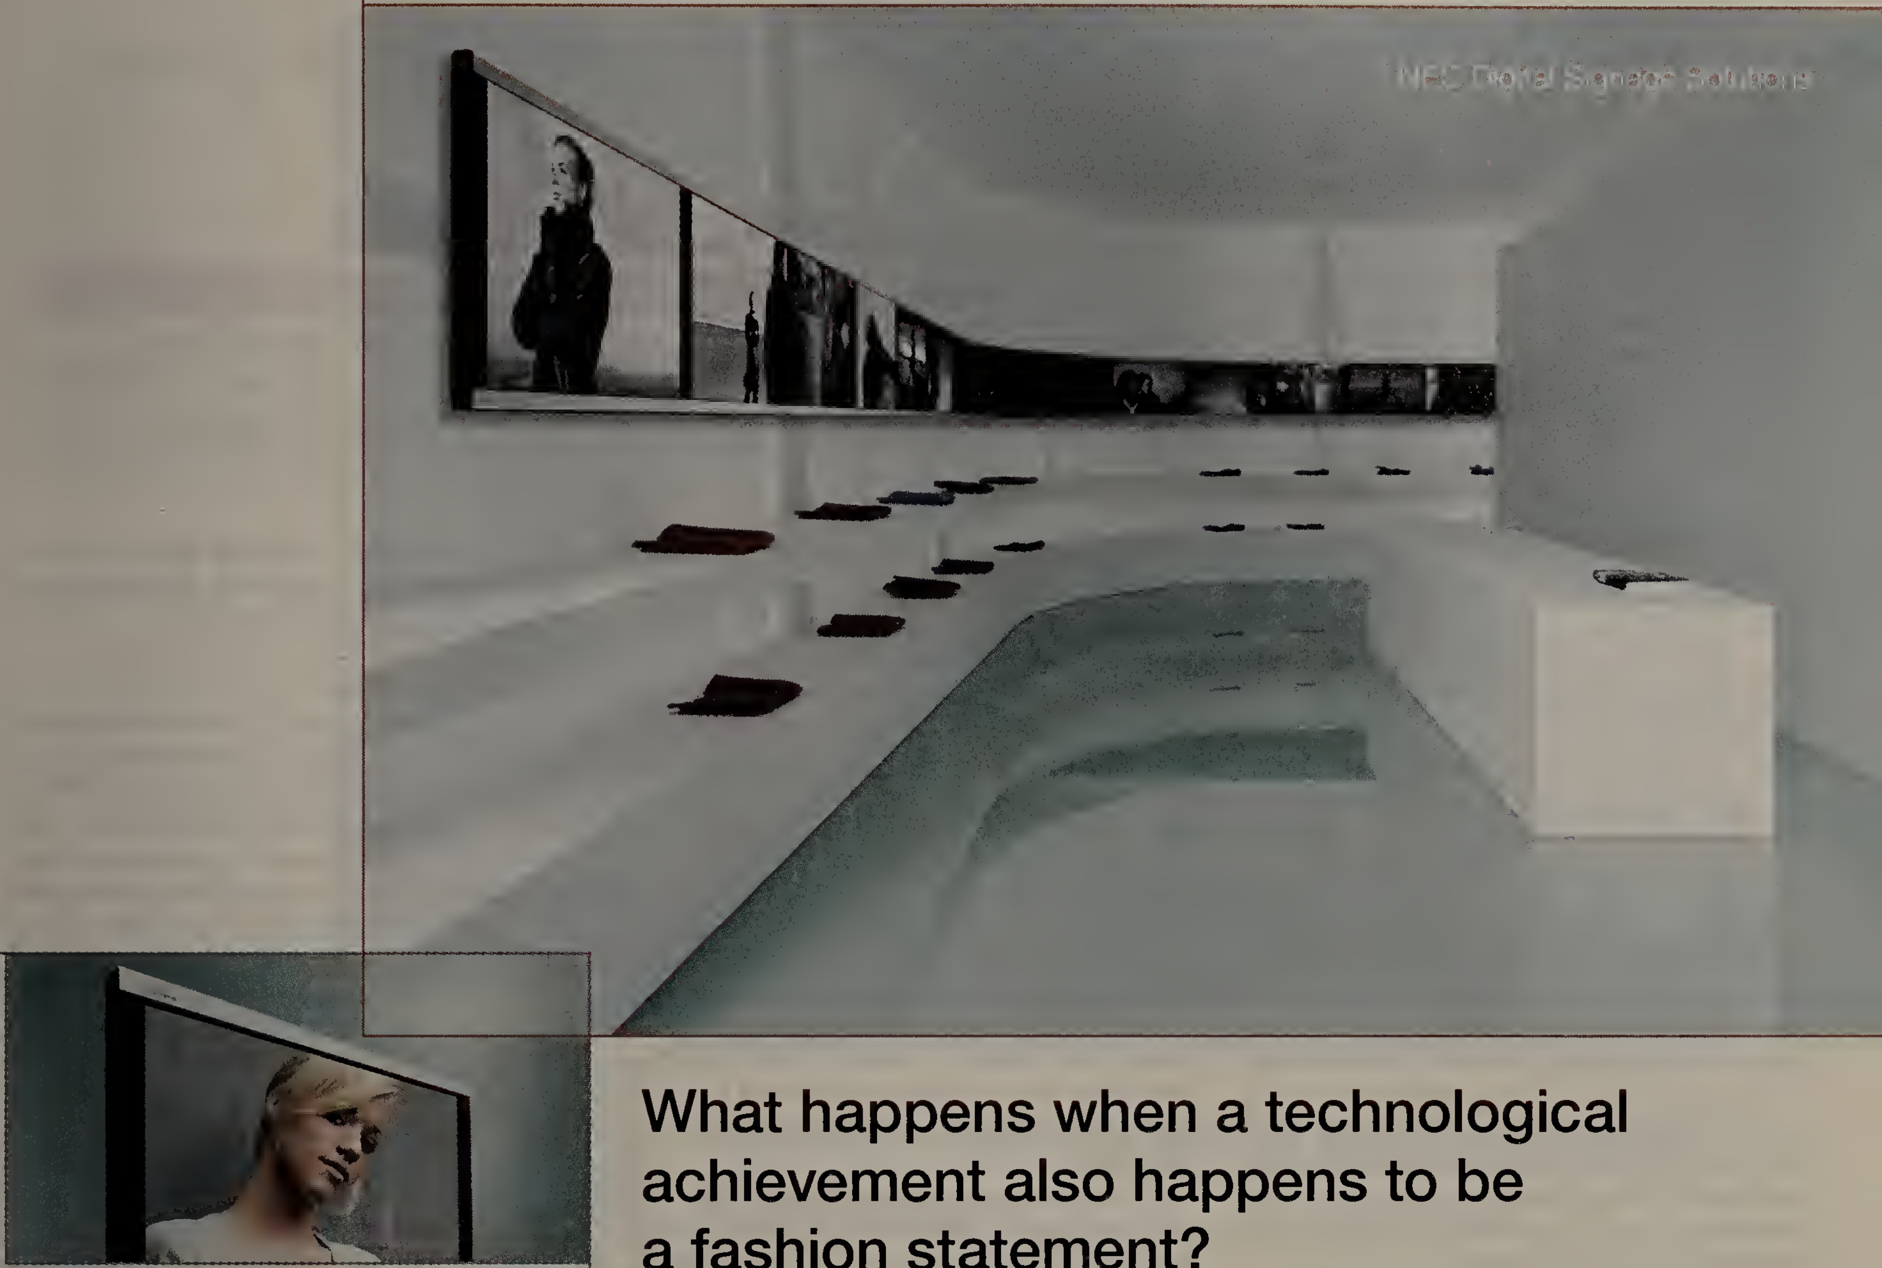

NEC Digital Signage Solutions

Pricing

continued from page 1

"It's tough from a software perspective to base pricing on margin, because the cost of goods and software is very low," Bament says. "We will look to see what the market value of our capabilities are. Obviously, being priced competitively is one of our value propositions, though we are not the cheap and cheerful low-end software solution. Typically, we will look at what the market bears."

Software prices can vary by country, a Microsoft spokesman notes. "As it relates to Windows, prices vary by region and are determined based on a variety of market-specific factors, including, but not limited to, exchange rate, local taxes, duties, local market conditions and retailer pricing decisions," the spokesman writes in an e-mail. "The primary principle in pricing Windows Vista was that comparable versions of Windows Vista would be priced the same as Windows XP."

Software prices are subject to negotiation, Bament says. Bluenote has guidelines for volume discounts "depending on the strategic nature of the customer," she says. Discounts are more likely if there are opportunities to deploy a product at a customer's subsidiaries, she adds.

With hardware, a vendor's cost of building products plays a much bigger role in pricing, says Bament, who has experience with hardware developers, such as Nortel and Motorola.

"You're pricing not just to market but taking into consideration the cost of the actual product," she says. "You'll find commodity hardware products, the margins are very small. With more customized and higher-end hardware, [the margins are] typically larger."

Figuring them out

How would you describe the typical pricing structures of leading enterprise network companies?

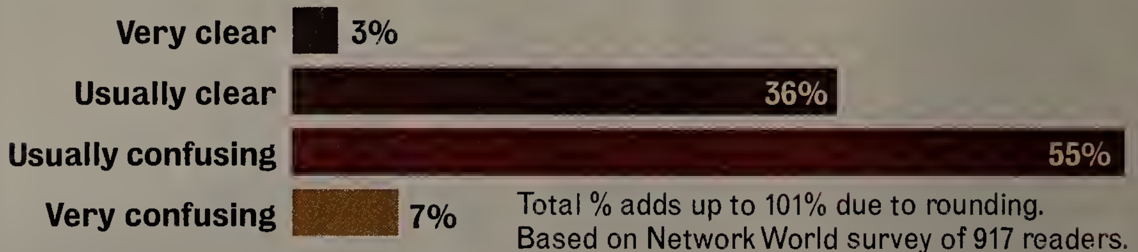

Total % adds up to 101% due to rounding. Based on Network World survey of 917 readers.

SOURCE: NETWORK WORLD TECHNOLOGY OPINION PANEL: WWW.NWWBETOPDOG.COM

For high-end computing systems, such as IBM Blue Gene supercomputers, pricing is still "heavily market-driven," says Herb Schultz, IBM's deep-computing marketing manager.

"The cost of things indicates a floor," he says. "It's not like we look at cost plus [some percentage]. You're always looking at market forces, competition, customer-buying behavior, what their capability to pay is. That's why IBM has other offerings, leasing and financing options. It's why we have Blue Gene in the on-demand center." IBM charges about $1.3 million per rack for the Blue Gene/P, its most advanced supercomputer, which was unveiled in June. The previous generation, the Blue Gene/L, also cost $1.3 million at one point, but IBM sales of the computer doubled this year after the company dropped its price to $800,000.

"Over time, some parts get lower in cost and we start getting economies of scale," Schultz says. "You want to maintain a price performance curve, which is always going down. In high-performance computing, the expectation is the price is always going down."

Art of price cutting

Cutting prices is often a good strategy, says Dan Clark, who worked in brand marketing at Digital Equipment in the mid- to late 1990s. A general manager thought Digital should raise the price of workstations to hit projected sales figures, Clark says. He argued that lower prices would increase sales volume enough to hit the dollar goal.

"I had pretty good evidence that there was elasticity in the pricing. The general manager thought pricing was inelastic," Clark says. "We won our argument by saying 'why don't we just sell one workstation for $1 billion?'"

After a price cut of 30%, Digital's workstation sales nearly tripled, moving the division from fifth to third in market share, Clark says.

Clark is now vice president of marketing at Lockdown Networks, which sells network access control appliances.

More than 60% of customers say that network-vendor pricing structures are confusing, Network World finds in a new poll (see graphic). Lockdown has tried a relatively simple pricing model, charging a flat fee of $25,000 for its appliance, while some appliance vendors will charge for the appliance itself, as well as per-user software costs, Clark says.

Lockdown's simple price structure is based on slightly more complex reasoning, however. The vendor figures customers usually pay between $25 and $100 to protect and maintain each desktop with antivirus software, Windows updates and patch management.

Lockdown's appliance targets enterprises with 500 to 1,000 users, so at $25,000 the per-user cost is typically between $25 and $50.

"It seems to be pretty on the mark for what people think the value is," Clark says. "We don't often get into intense negotiation for cost per user."

The other major pricing issue is maintenance and support. Clark says Lockdown typically charges 25% of product cost for support, more than the industry's 15% to 18%. Lockdown says that it offers better-than-average value because it provides updates twice a day. ■

So, what is TCO?

When calculating total cost of ownership, the price you pay vendors for IT products is just the tip of the iceberg. Per-user TCO is about 4.5 times higher than the actual price of hardware and software in typical scenarios, when users who provide informal IT support, administration, downtime and operation costs are factored in. That's according to research issued last November by Gartner.

TCO is a mix of direct and indirect costs related to assets and tasks. Nearly half of a typical TCO is from users who perform informal technical support, perhaps because of an IT staff shortfall, Gartner said in another report in February.

Gartner says these "end-user operations costs" tend to be hidden, unbudgeted and poorly accounted for. Labor costs can be reduced, however, by making strategic investments in operations assets, such as help desk automation, systems management tools and updated operating systems.

A thorough analysis of these factors can help an IT department build the business case for new products and upgrades.

"Infrastructure and operations funding is hard to justify and obtain. The dynamics of TCO can dramatically improve the business case for such investments when indirect costs are considered," Gartner's Lars Mieritz and Bill Kirwin write.

Gartner defines total cost of ownership as the "holistic view of costs across enterprise boundaries over time." The definition has changed over the years to include non-IT costs that can be related to IT, such as human resources and facilities.

NEC Next-Generation Thin Client Solutions

This CTO knows vendor weaknesses

You already know IT prices are too high. Here's how to make vendors agree.

BY JON BRODKIN

Finding a bargain

Where do you turn first for a bargain on enterprise network products?

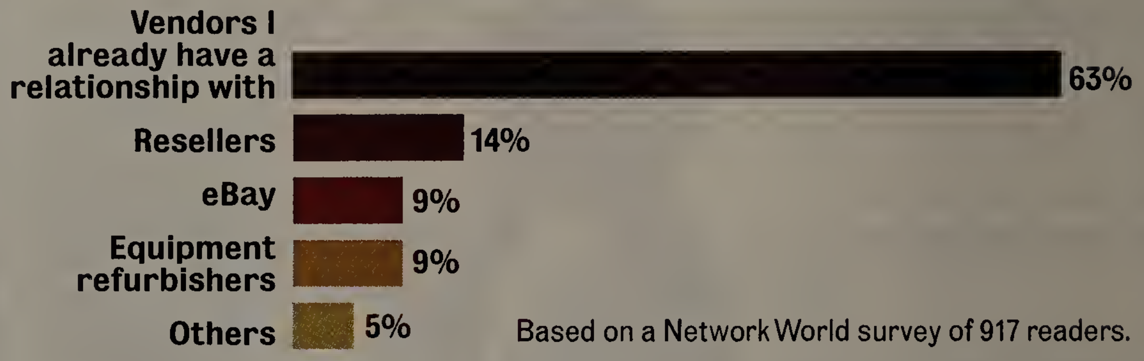

Based on a Network World survey of 917 readers.

SOURCE: NETWORK WORLD TECHNOLOGY OPINION PANEL: WWW.NWWBETOPDOG.COM

Everybody wants a bargain. But when it comes to the complex world of IT products, finding a deal or even knowing what something should cost can be tricky. Sixty-two percent of IT buyers say the pricing structures used by enterprise network vendors are "usually confusing" or "very confusing," a new Network World poll finds.

Infocrossing CTO Dave Leonard, however, has figured it all out. With extensive experience as a buyer and seller, Leonard knows all the tricks for getting discounts from vendors.

There are basically three things that motivate a typical vendor, Leonard says. The first is obvious: The vendor want new sales. The second is less obvious: The vendor aims to displace competitors, and may even set aside "displacement funds" specifically to give discounts to customers who agree to get rid of a competitor's product and replace it with the vendor's. The third driver is the fear of losing ongoing revenue streams from maintenance and support costs.

Customers can use this knowledge to get better deals, even in noncompetitive markets. Salespeople want to close deals before the end of the quarter because they are under constant pressure to meet goals for each three-month period. In other words, make a sales representative sweat for a few extra weeks toward the end of a quarter and you might get a discount.

"Even if there's not a competitive situation, using time against the vendor gives them the opportunity to sweeten the deal," Leonard says.

Leonard pulled out all the stops recently when Infocrossing, an IT outsourcing provider based in New Jersey, embarked on a standardization initiative across four data centers. The project saved the company $14 million through consolidation of labor, software and other costs.

Measuring up management software

Infocrossing, which operates in 12 states, has quadrupled in size through three major acquisitions over the past four years. The company ended up with data centers running three different server management tools, from CA, IBM's Tivoli division and NetIQ. Infocrossing decided to standardize on NetIQ after an evaluation of the products but didn't tell the three vendors that the decision had already been made. The first step was to build a business case showing each vendor how much it would cost internally to use its products.

"We did a complete economic analysis to get to 'what will it take us to get to a single platform,'" Leonard says. "The cost is kind of what we presented back to the vendors. . . . The idea behind that was to get them to understand that our cost of using their product was far greater than the actual cost the product was going to be."

Leonard told the vendors that Infocrossing didn't want to be flooded with consultants, because its own employees would have to run the system. "We did say you can help the overall economic case by affecting how much maintenance we pay on our existing install base," Leonard says.

He also asked vendors to loosen restrictions on existing contracts, such as clauses that prevent a product license from being used in more than one data center.

Each time a vendor offered a proposal to entice Infocrossing, the company was able to bounce the idea off the other two vendors and ask them to do better. "We're trying to end up with something that's defensible on both sides, because they have to sell it internally," he says.

After a negotiation period of three months, Infocrossing got a deal from NetIQ that Leonard says will save the company "seven digits" over the next five years.

The company was already running NetIQ on about 1,500 servers and wanted to standardize across 5,000. After the wheeling and dealing, the license charges for the additional 3,500 servers were "negligible" because NetIQ funded the cost with competitive displacement money.

"It went from a significant to an insignificant cost of the whole operation," Leonard says.

Infocrossing also focused on maintenance costs, because paying 20% of list price, as vendors would prefer, "can just kill you," he says.

Leonard's goal is typically to pay 20% of the acquisition price and negotiate clauses that limit price increases related to future acquisition of licenses. Immature IT buyers often make the mistake of focusing only on upfront costs, when future costs for maintenance and additional license acquisitions can turn a seemingly good deal into a bad one, he says.

The process Leonard used to negotiate lower server management costs was replicated across 20 or 30 products in the standardization initiative, making software a significant portion of the cost savings achieved in the whole project.

Infocrossing still uses many IBM and CA products in areas other than server management. Infocrossing ended up paying NetIQ more overall than it did previously, but the cost per server is "way less," Leonard says. He can't say exactly how much it paid due to a nondisclosure agreement.

"If they don't hamstring you with a nondisclosure agreement, that generally means you didn't get good pricing," Leonard says. "They don't want our pricing available to the general public."

Infocrossing often finds itself on the other side of the table, when its own customers ask for discounts. "Sometimes, we'll say 'absolutely,'" because when Infocrossing's hardware costs go down, it makes sense to pass some savings on to customers, Leonard says.

"If there's not a basis where our costs have gone down, we go back to them and say 'hey, here's what our costs are, there isn't anything that's changed since we did the deal before. It was a good deal then and it's still a good deal now,'" Leonard says.

Leonard says a vendor that is logical and unemotional can typically convince a customer that the price is right, even if the customer has asked for a discount.

"The customers don't know, largely, how much things should cost," he says. "The more confidence they get from us that we know how much things should cost," the more confidence they will have in the pricing. ■

Go online and get your FREE efficiency rating today!

Or fill out this card for the following white papers:

Efficiency™ Quotient
DATA CENTER

□ White Paper #114 **"Implementing Energy Efficient Data Centers"**
□ White Paper #63 **"AC vs DC Power Distribution for Data Centers"**
□ White Paper #113 **"Electrical Efficiency Modeling for Data Centers"**

□ **YES!** Please send me my FREE white papers. □ **NO,** I'm not interested at this time, but please add me to your mailing list.

Name: ______________________ Title: ______________________

Company: ______________________

Address: ______________________ Address 2: ______________________

City/Town: ______________ State: __________ Zip: __________ Country: __________

Phone: ______________ Fax: ______________ e-mail: ______________

□ ***Yes!*** **Send me more information via e-mail and sign me up for APC PowerNews® e-mail newsletter.** **Key Code: x420x**

What type of availability solution do you need?
□ UPS: 0-16KVA (Single-phase) □ UPS:10-80kVA (3-phase AC) □ UPS:80+ kVA(3-phase AC) □ DC Power
□ Network Enclosures and Racks □ Precision Air Conditioning □ Monitoring and Management
□ Cables/Wires □ Mobile Protection □ Surge Protection □ UPS Upgrade □ Don't know

Purchase timeframe? □ <1 Month □ 1-3 Months □ 3-12 Months □ 1 Yr. Plus

You are (check 1): □ Home/Home Office □ Business (<1000 employees) □ Large Corp. (>1000 employees)
□ Gov't, Education, Public Org. □ APC Sellers & Partners

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 36 WEST KINGSTON RI

POSTAGE WILL BE PAID BY ADDRESSEE

APC

ATTENTION CRC:x420x
132 FAIRGROUNDS ROAD
PO BOX 278
WEST KINGSTON RI 02892-9920

2
4
1
3

Get up to scratch on metal whiskers

Troublesome filaments can wreak havoc in data centers

BY RYAN DEBEASI

Depending on whom you ask, the data-center phenomenon of metal whiskering is either a relatively uncommon fluke or a crisis waiting to happen.

Whiskering is caused either by stress from a particular manufacturing technique used by makers of servers, floor tiles and other products, or by a cornucopia of factors. Some say the problem can be avoided by not using old or inexpensive materials, while others say new research is required to eliminate the threat.

Most data-center equipment manufacturers are taking measures to prevent metal whiskers — troublesome, tiny filaments that can form on their products' zinc and tin coatings

Still, data centers with old or inexpensive materials or equipment run the risk of whiskers forming, breaking off, getting into computers and short-circuiting them. Metal whiskers may cause unusual, sporadic problems, or they may cause a data center's power supplies to short out en masse. Since *Network World* first covered metal whiskers in 2004 (www.nwdocfinder.com/1422), new research and environmental legislation have changed how people approach the issue.

Metal whiskers get a close-up

Zinc whiskers shown growing on the underside of an old floor tile (inset) and a closer look (below) using a powerful microsope.

Go online to view a slideshow of more metal whisker pictures. **www.nwdocfinder.com/1421**

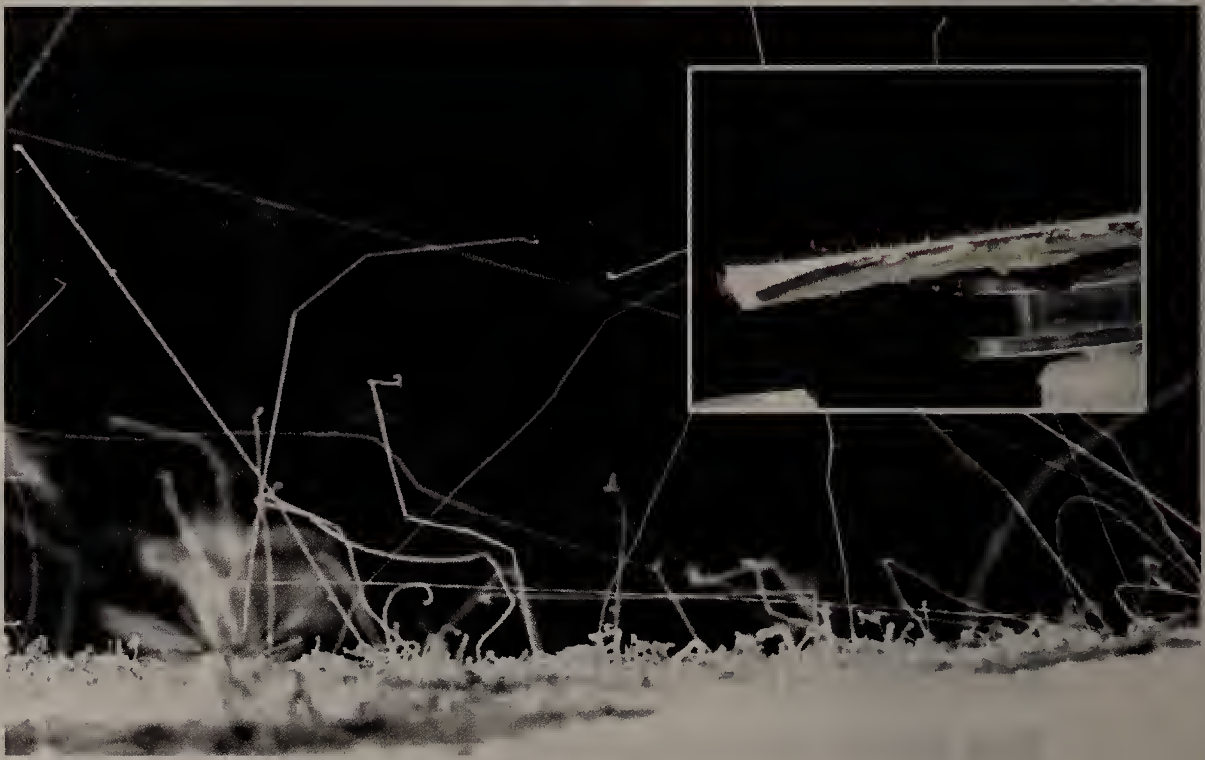

PHOTOS COURTESY OF THE NASA ELECTRONIC PARTS AND PACKAGING PROGRAM

Explaining metal whiskers

The source of metal whiskers is steel that has been electroplated with zinc or tin to prevent rusting, according to Robert Sullivan, senior consultant at the Uptime Institute in Santa Fe, N.M. When manufacturers deposit zinc or tin through electroplating, he says, the process can introduce stresses that cause whiskering. Zinc- and tin-plated metal has been used in computers, server racks, floor tiles and the like.

Sullivan says that hot-dipped galvanizing — a process in which metal is dipped into molten zinc or tin — does not produce whiskers. He adds that if metal is electroplated, a powder coating process can be used to prevent whiskers from forming.

According to Sullivan, most manufacturers of metal tiles and the like use metals and processes that don't cause whiskering, but old or inexpensive materials may pose a risk. He says that in one data center that was less than five years old, he found metal whiskers on an economy-grade metal bar that was used to support ceiling tiles. Only the bottom of the bar was powder-coated, and zinc whiskers formed on the other sides. Sullivan says metal whiskers generally take about two years to form, although he has seen them crop up in as little as six months.

Mixing lead into tin or zinc prevents metal whiskers from forming, and new, lead-free solder could introduce a new source of whiskers. "I think we're just starting to see the tip of the iceberg on that," says Rich Hill, president of data-center cleaning company Data Clean. Sullivan disagrees: "I don't see that soldering is an exposure to the creation of either zinc or tin whiskers," he says.

Who has metal whiskers?

Documentation of metal-whisker problems in data centers is hard to come by.

"Whiskering is something that people keep close to their vest," Hill says. "You don't want your clients to know you have [metal] whiskers." He adds that this especially is true in the case of collocation facilities and other organizations whose reputations are built on data-center reliability. Hill said that although he didn't know of any clients that would be willing to talk about metal whiskers, he has seen the problems the tiny filaments can cause. "We've heard things go 'pop'; we've had clients that lost a hundred power supplies in a weekend" due to metal whiskers, he said.

NASA is no stranger to metal-whisker problems: the organization runs a Web site (see www.nwdocfinder.com/1423) that covers metal-whisker research, and metal whiskers have caused failures in NASA equipment, including a flight control system. According to a 2006 NASA presentation, metal whiskers have caused equipment failures in satellites, telecommunications equipment, missile programs and nuclear power plants. A presentation by the U.S. Department of Energy Office of Environment, Safety and Health Evaluations says metal whiskers caused a nuclear power reactor to shut down in April 2005.

Layne Maly, director of communications for data-center user group AFCOM, found less evidence of problems with metal whiskers. She said in an e-mail that she had asked the majority of the group's members whether they had encountered such problems: "The responses I've received back all say the same thing — 'I have not had a problem with metal whiskers, and I don't know anyone who has.'"

Stress from manufacturing might not be the only factor in the growth of metal whiskers. Research on the effects of humidity, electrical current and other factors is contradictory and inconclusive, according to NASA's metal-whisker Web site, and so the organization argues these factors should not be ruled out. In addition, NASA says that stress from sources other than manufacturing — for example, scratching or bending metal — also could cause whiskers to grow. For his part, Data Clean's Hill says there are "no conclusions out there as to what causes whiskering." The Uptime Institute's Sullivan acknowledges that higher temperatures can speed the process, but he says that stress is the root cause of the phenomenon and that humidity and other factors have no effect on whiskers. ■

Cisco playing network defense

Its Self-Defending Network strategy moves forward via IronPort buyout

BY JIM DUFFY

Cisco's 6-year-old Self-Defending Network strategy for securing converged networks remains a work in progress: Acquisitions and internal developments are moving it forward even as customers push Cisco to go above and beyond its initial plans.

Cisco spends $400 million annually — roughly 10% of its total R&D budget — on security. The company's aim with SDN is to integrate security into all aspects of a converged data, voice and video network with a focus on secure connectivity, threat defense, and trust and identity management.

In June, Cisco provided its most recent update on SDN after its acquisition of IronPort Systems, a privately held developer of e-mail and Web security products. Cisco said IronPort ushered in Version 3.0 of SDN (Version 1.0 involved Cisco's recognition that security is more than point products; Version 2.0 comprised building those capabilities into Cisco products.)

Cisco plans to port IronPort's SenderBase reputation services onto Cisco Adaptive Security Appliance firewalls by the first half of 2008. It also plans to port SenderBase to other key security or routing platforms, such as the Integrated Services Routers and Mitigation Analysis and Response System. Integration with Cisco and third-party network admission control (NAC) products also is expected.

"If they can now get e-mail security, Web security — basically all the secure messaging technologies — into that mix they've got a bigger story," says Charlotte Dunlap, senior analyst at Current Analysis.

Dunlap is keeping an eye on how Cisco might take advantage of an existing relationship between IronPort and Vontu, a developer of software that analyzes content and authorizes user access at endpoints to protect against data leakage.

"I'd really like to hear their data-leakage story," says Dunlap, who compares Cisco's purchase of IronPort to Secure Computing's acquisition of CipherTrust last year. "[IronPort does not offer] the level of depth that the data-leakage prevention providers do."

Cisco intends to maintain IronPort's ties to Vontu and exploit the relationship for inclusion in the SDN architecture, according to Jeff Platon, vice president of security marketing at Cisco.

"I think of that as a part of the solution but I do see a variety of other parts of the portfolio that are being enhanced to participate in a more comprehensive data-leakage solution," Platon says. "It's a tough problem — you can't just rely on one methodology."

Cisco's earlier buyout of FineGround in May 2005 fits into the plan. Pieces of the FineGround technology have found their way into the Application Control Engine blade for Cisco Catalyst 6500 switches, Platon says. ACE is a key component of SDN's data-center security component, in which application connection requests to server farms are inspected for legitimacy and outbound content authorization, and filtered for malware.

Beyond SDN 3.0, Cisco plans to build greater collaboration among network-, content- and application-layer services, Platon says. Reputation services will broaden to include users, perhaps through what Platon calls a global passport service yet to be created by a public- or private-sector enterprise.

"You're going to have to have some way to determine [who someone is] with some semblance of accuracy," Platon says. "Reputation on a user basis is one of the possibilities that I think has a great amount of likelihood to come to pass."

That hits home with Pacific Gas & Electric, which is undergoing a business transformation whereby it is constructing centralized resource management centers. PG&E relies on Cisco for desktop-to-core connectivity, process and security requirements, says Paul Nielsen, supervisor of LAN/WAN services at the utility. "At those centers, where we've deployed the majority of their products, is where to build a self-defending network."

PG&E uses Cisco PIX firewalls, Clean Access NAC appliances, VPN concentrators and Firewall Service Modules on the Catalyst 6500 LAN switches, as well as the "latest and greatest" security features on Cisco switches and routers, Nielsen says. "It's more than just deciding if you have the right certificate or the right credentials; it's more important that we find out if there's a watermark on the PC, if this is a PG&E person," he says.

An announcement last week by Cisco and Intel might help. Intel enhanced its vPro processor technology with a Cisco-certified embedded trust agent that offers Cisco customers the ability to manage systems without lowering the security on IEEE 802.1x networks and Cisco SDN products. Nielsen says PG&E hasn't been briefed yet on Cisco's road map for that. But where SDN currently fits is in spots where PG&E is installing new Cisco infrastructure.

"Where we've had problems is where we have legacy systems," Nielsen says. "If a company buys into the Cisco solution and they buy all of the pieces, it works great; but you've got to have all of the pieces there. You can't do clean access NAC on a Catalyst 1900 switch that was built six or 10 years ago; it just doesn't work."

Nielsen notes that this issue is industry-wide, not Cisco-specific.

Customer, analyst wish lists

PG&E would like to see Cisco take SDN into the realm of virtualization, especially with intrusion detection.

See Cisco, page 52

Key elements of Cisco's Self-Defending Network strategy

- **Cisco Security Agent** — Desktop and server agent software for prevention of malware intrusions.
- **IronPort SenderBase** — Reputation services for Web-server and e-mail security.
- **Reactivity** — February 2007 acquisition for XML gateway and security hardware.
- **FineGround** — Acquired in May 2005 for bandwidth optimization appliances designed to accelerate, secure and monitor application delivery in the data center.
- **Vontu** — Data leakage prevention company has partnered with IronPort/Cisco.

Turn back network time.

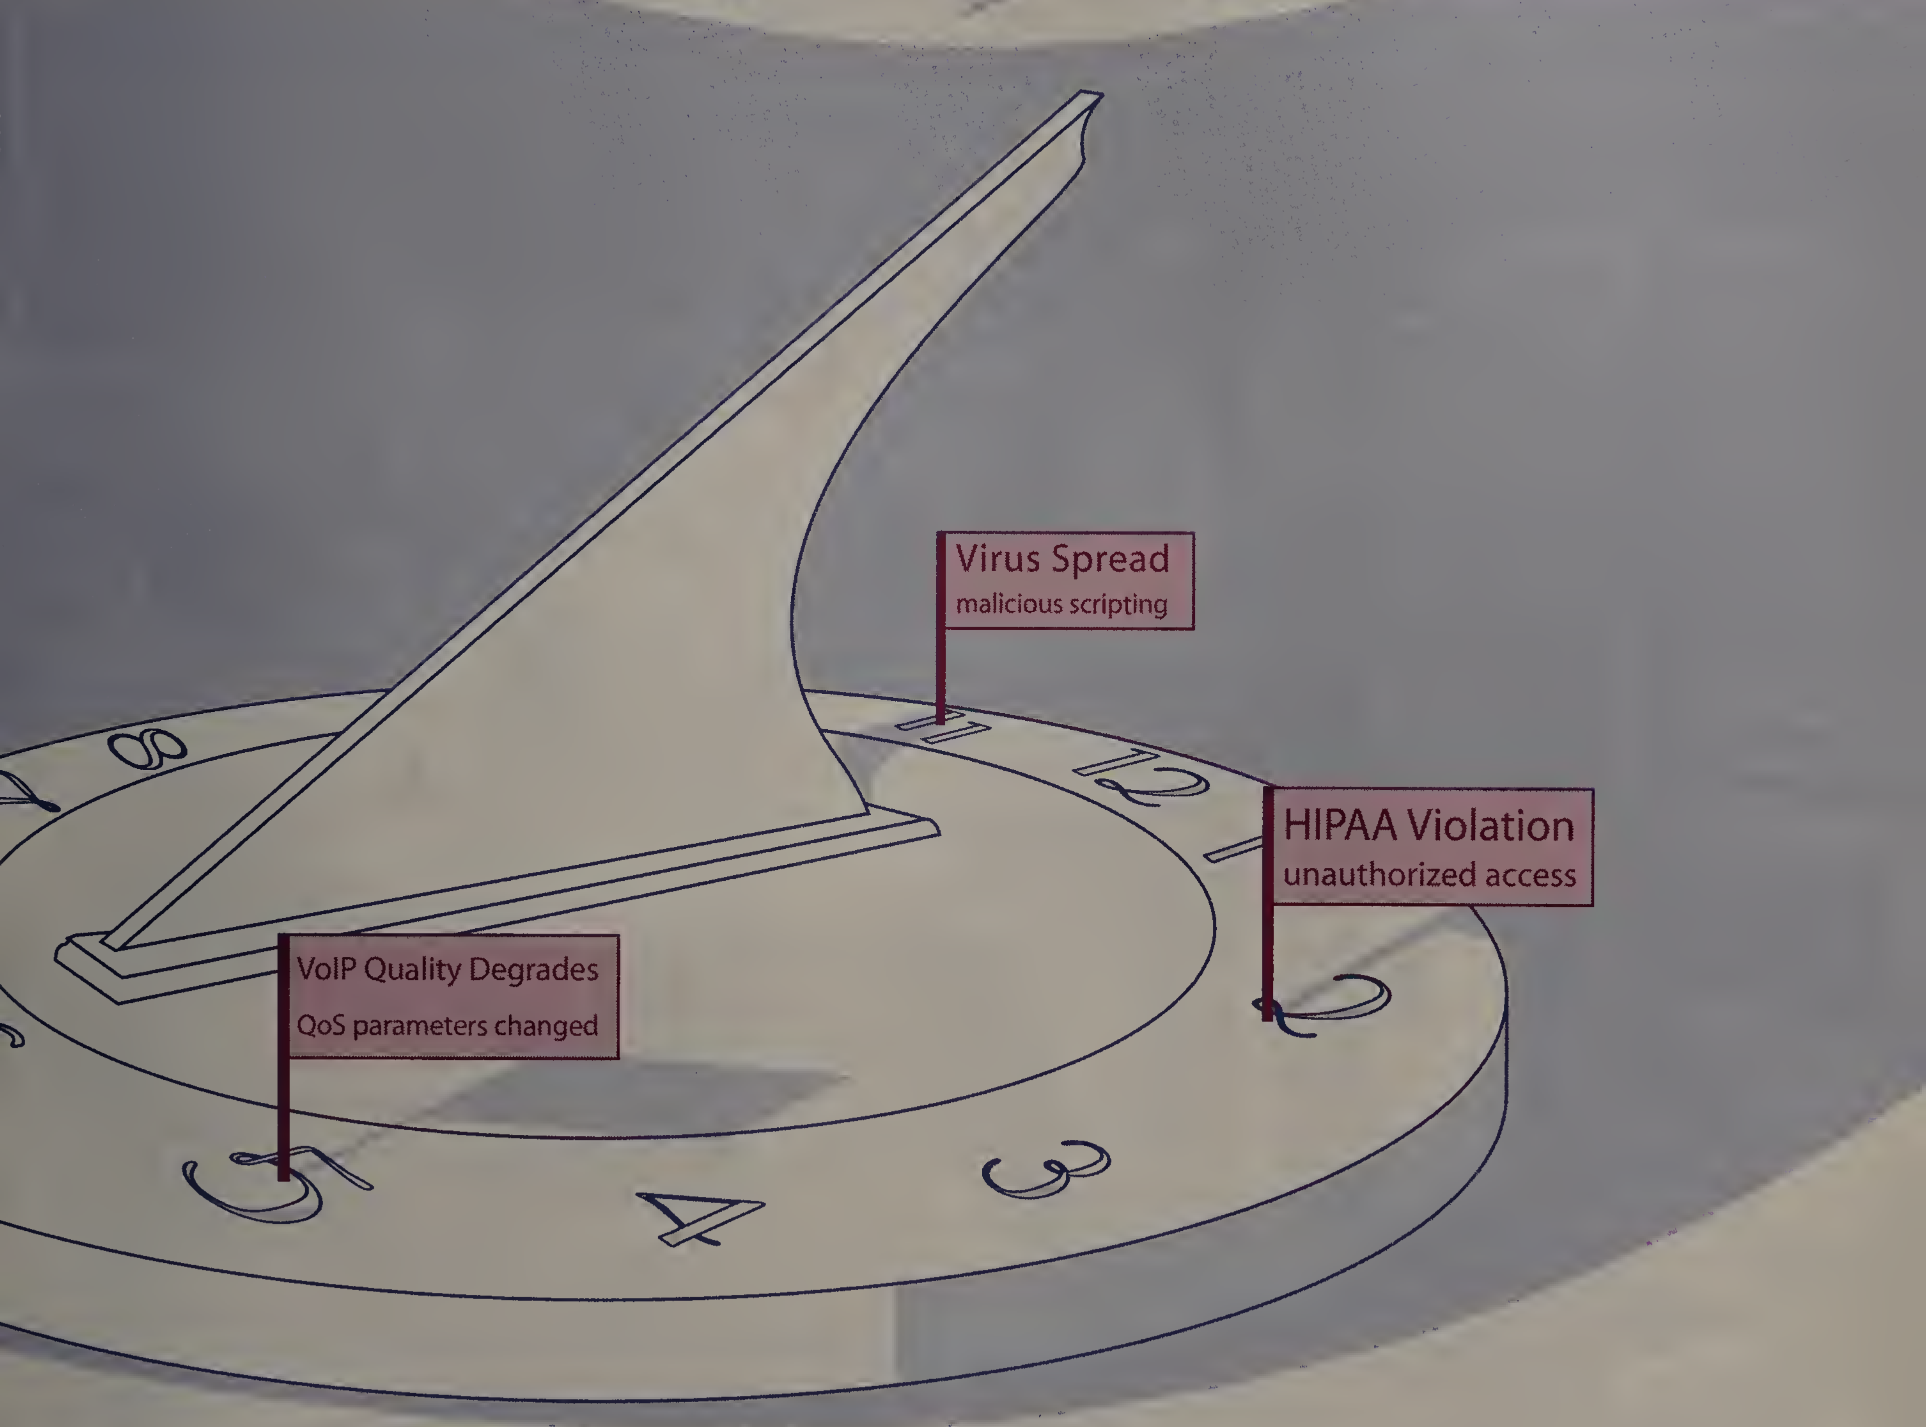

Stop missing critical events.

For a trusted approach to problem resolution rely on the Network Instruments® GigaStor™ appliance. Everything is recorded—every packet, every protocol, every transaction for hours, days, even weeks. The unique GigaStor interface provides an effective way to go back in time to determine not only when the application went down but *why*.

Resolve intermittent problems, track compliance efforts, isolate VoIP quality issues, and more on the most complex WAN, Gigabit, and 10 GbE networks. Find out how you can turn back the clock with the GigaStor. After all, your network history shouldn't be a thing of the past.

GigaStor: Get proof. Take action. Move forward.

NETWORK® INSTRUMENTS

Learn more about GigaStor. 800-526-5958
www.NetworkInstruments.com/TimeTravel

IBM

FRAGILE
HANDLE WITH CARE
Tivoli.

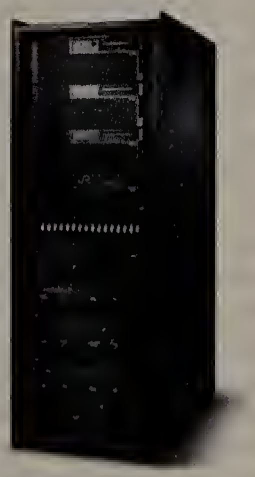

Understanding federated identity

BY WILLIAM STALLINGS

Federated identity management is a relatively new concept that is an extension of identity management, which is a centralized, automated approach to regulating access to enterprise resources by employees and other authorized individuals.

The focus of identity management is defining an identity for each user (human or process), associating attributes with the identity and enforcing a means by which a user can verify identity. Once implemented, identity-management systems support single sign-on (SSO), the ability of a user to access all network resources after a single authentication.

Federated identity management refers to the agreements, standards and technologies that enable the portability of identities, identity attributes and entitlements across multiple enterprises and numerous applications, supporting thousands, even millions, of users.

When multiple organizations implement interoperable federated identity schemes, an employee in one organization can use SSO to access services across the federation with trust relationships associated with the identity.

Beyond SSO, federated identity management provides other capabilities. One is a standardized means of representing attributes. Increasingly, digital identities incorporate attributes other than an identifier and authentication information (such as passwords and biometric information). Attributes can include account numbers, organizational roles, physical location and file ownership. And a user may have multiple identifiers associated with multiple roles, each with its own access permissions.

Another key function of federated identity management is identity mapping. Security domains may represent identities and attributes differently. Further, the amount of information associated with an individual in one domain may be more than is necessary in another domain. The federated identity-management protocols map identities and attributes of a user in one domain to the requirements of another domain.

A generic federated identity-management architecture (see graphic) includes identity providers and service providers. The identity provider acquires attribute information through dialog and protocol exchanges with users and administrators.

Service providers are entities that obtain and employ data maintained and provided by identity providers, often to support authorization decisions and to collect audit information. For example, a database server or file server is a data consumer that needs a client's credentials to know what access to provide to that client. A service provider can be in the same domain as the user and the identity provider or in a different domain.

The goal is to share digital identities so a user can be authenticated once and access applications and resources across multiple domains. The cooperating organizations form a federation based on agreed-upon standards and mutual levels of trust.

Federated identity management uses a number of standards as the building blocks for secure identity exchange. In essence, organizations issue some form of security tickets for their users that can be processed by cooperating partners. Identity federation standards are thus concerned with defining these tickets, in terms of content and format, providing protocols for exchanging them and performing a number of management tasks. These tasks include configuring systems to perform attribute transfers and identity mapping, and performing logging and auditing functions.

The principal standard for federated identity is the Security Assertion Markup Language (SAML), which defines the exchange of security information between online business partners.

SAML is part of a broader collection of standards being issued by the Organization for the Advancement of Structured Information Standards for federated identity management. For example, WS-Federation enables browser-based federation; it relies on a security token service to broker trust of identities, attributes and authentication between participating Web services.

The challenge with federated identity management is to integrate multiple technologies, standards and services to provide a secure, user-friendly utility. The key is the reliance on a few mature standards widely accepted by industry. Federated identity management seems to have reached this level of maturity.

Stallings is coauthor of the new book, Computer Security: Principles and Practice. *Contact him at ws@shore.net.*

How federated identity works

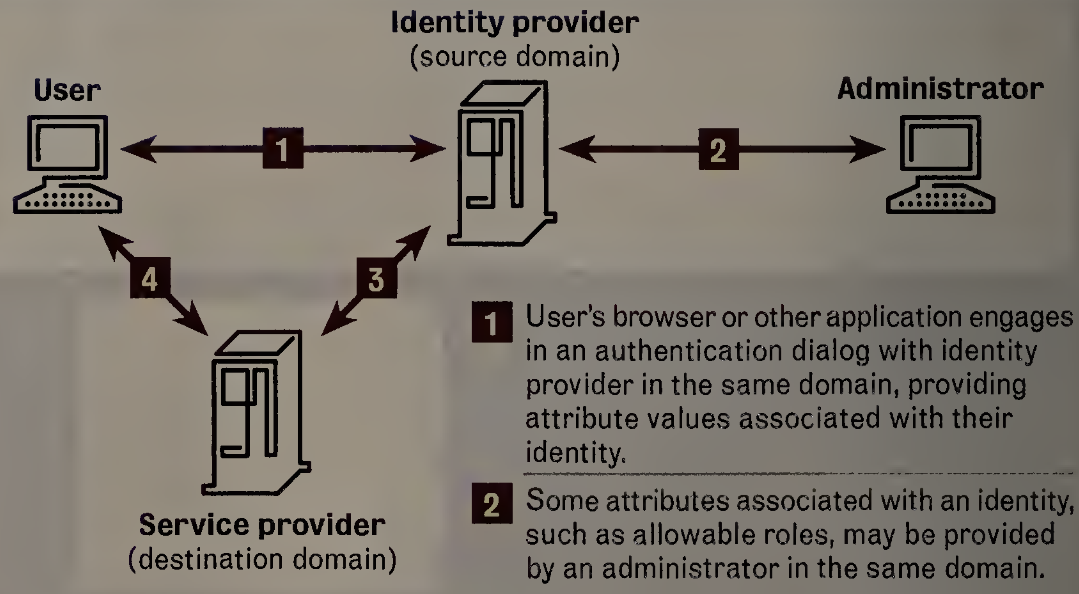

1 User's browser or other application engages in an authentication dialog with identity provider in the same domain, providing attribute values associated with their identity.

2 Some attributes associated with an identity, such as allowable roles, may be provided by an administrator in the same domain.

3 A service provider in a remote domain that the user wants to access obtains identity information, authentication information and associated attributes from the identity provider in the source domain.

4 Service provider opens session with remote user and enforces access control restrictions based on user's identity and attributes.

GEARHEAD

Mark Gibbs

Feedback: Outlook to iCal

This week I have a lot of reader feedback to deal with. First up is a recommendation from longtime reader Gar Nelson regarding my recent quest (www.nwdocfinder.com/1431) for a program that would catalog and manage my collection of data CDs.

Gar suggested Readerware, and while it doesn't specifically address cataloging data disks, it can catalog all sorts of media, and runs on Windows, Mac, Linux and Palm.

What amused me was when I went to check out the product on the Readerwear site I found an offer (www.nwdocfinder.com/1432) that is a blast from the past: A free CueCat!

Those of you who didn't fry your brains during the Internet Bubble may remember this oddly packaged device. The CueCat is a USB- or PS/2-based bar-code reader developed by the now-defunct Digital Convergence Corporation in the shape of a stylized cat.

Starting in 2000 Digital Convergence distributed tens of thousands of these devices to consumers in concert with a huge marketing campaign to get magazines and retailers to put bar codes in printed material that the Cat could read. I don't have space to go into the flawed business model, the hacking of the CueCat devices, and the mass exposure of consumer data (see the Wikipedia entry at www.nwdocfinder.com/1433) but the result, unsurprisingly, was that Digital Convergence went belly up in 2005.

The CueCat is indeed free, but only with the purchase of one of the Readerware bundles, which start at $85. I have yet to take a serious look at this software, so if you have, let me know your thoughts.

Of all the Gearhead columns from the past year, my recent column on trying to automate the export of calendar data from Outlook to iCal generated the most mail. From this I would guess that things to do with Outlook and Exchange feature very heavily in your lives — perhaps more so than I would have guessed. So, tell me: What are your biggest issues with Outlook? What problems are you trying to solve?

My problem seems to have a number of solutions. Reader Bruce Gerson suggested using a program called GroupCal from Snerdware (www.nwdocfinder.com/1434), but you have to be also using Microsoft Exchange, which I'm not. GroupCal looks very promising and even provides iPhone support, albeit with limitations. At $55 for a single seat, GroupCal looks like a steal, and Bruce says it is extremely easy to set up.

What are your biggest issues with Outlook?

Other suggestions involved external tools to drive Outlook through its user interface, but what I wanted was something that would work using Outlook scripting and/or APIs.

Reader Joe Kendal wrote that he took the outlook2ical code written by Norm Jones that I mentioned last week and converted it to Visual Basic to run outside of Outlook. Using the Redemption DLL to get around Outlook security, Joe created what he describes as a working solution, except he thinks it is "not production worthy ... There needs to be some extra testing and error handling for it to be 100%."

I haven't had a chance to test Joe's code but he says that anyone who wants a copy is welcome. Drop a message to gearhead@gibbs.com with the subject Kendal and the code will automatically be sent to you, then tell me what you think.

Gibbs will always read feedback sent to gearhead@gibbs.com.

Keith Shaw

COOLTOOLS

Not much advantage with the 7501

The scoop: HTC Advantage 7501, by HTC America, about $900, plus wireless service.

What it is: Somewhere between a PDA, smartphone, ultra-mobile PC and an iPhone lies the HTC Advantage 7501. The Windows Mobile 6-enabled device combines several features for mobile professionals into a palm-sized device, including a mobile phone, PDA, digital camera, wireless e-mail device, mobile multimedia player and a GPS navigation system.

The device is designed for professionals who want to ditch a notebook and carry around something smaller but still be able to access heavy-duty business applications, including Microsoft Office and e-mail. The Advantage includes a 5-inch touchscreen, an 8GB hard drive (with additional microSD memory card support), a 3-megapixel digital camera, and Direct Push support for synchronizing with Microsoft Exchange. Network connectivity options include 3G wireless WAN (it supports the HSDPA network — we tested ours on the AT&T wireless network), built-in Wi-Fi and Bluetooth. The GPS includes the TeleNav application for driving directions and general navigation.

Why it's cool: A magnetic attachable keyboard makes this device more fun than a normal stodgy Windows Mobile device. The keyboard enables for text input that's easier than using an on-screen keyboard or trying to deal with handwriting recognition with a stylus. The 3-megapixel digital camera application makes for some of the best-looking photos I've seen from a mobile phone-type device.

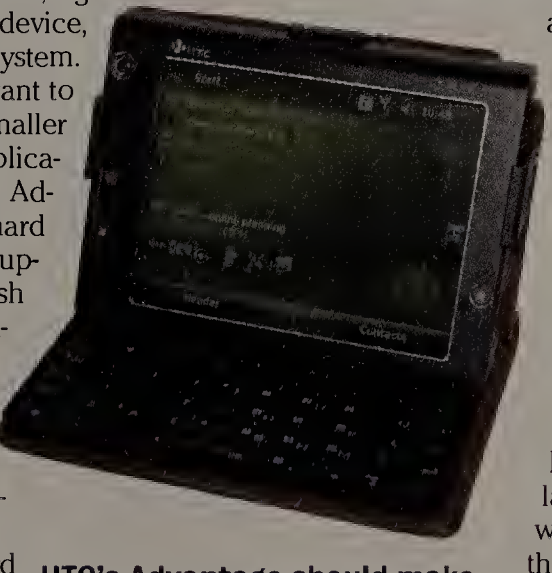

HTC's Advantage should make Microsoft fans happy.

The GPS application from TeleNav was very good, as it could take advantage of real-time traffic data through the wireless LAN connection.

Some caveats: Trying to connect the device to my Exchange server for e-mail access and synchronization was an exercise in frustration. An incompatible VPN connection prevented me from any wireless synchronization. The only way I could connect to Outlook was through a USB cable connected to my notebook. This meant that I could only offload e-mail to the device, I wouldn't be able to instantly respond as I could with a BlackBerry, for example.

Even the USB connection was tricky — after several nonconnection errors, I had to wade through the 250-page manual to discover that a check box in the "USB-to-PC" applet on the Advantage needed to be unchecked in order to force the device into a serial USB setting, rather than the "advanced network mode."

Although the digital camera took great photos, it had a klunky interface with lots of icons that were hard to decipher, making it more difficult to figure out if I had the right setting selected.

Bottom line: If you have a high tolerance for Exchange, ActiveSync and Windows Mobile installation procedures, this device might interest you. But with the vast number of other mobile devices out there (this device is more expensive than an iPhone), there's not much advantage in this Advantage.

Grade: 3 stars (out of five).

Shaw can be reached at kshaw@nww.com. New Cool Tools video show every Thursday, and Twisted Pair podcast every Friday at www.networkworld.com.

IBM.

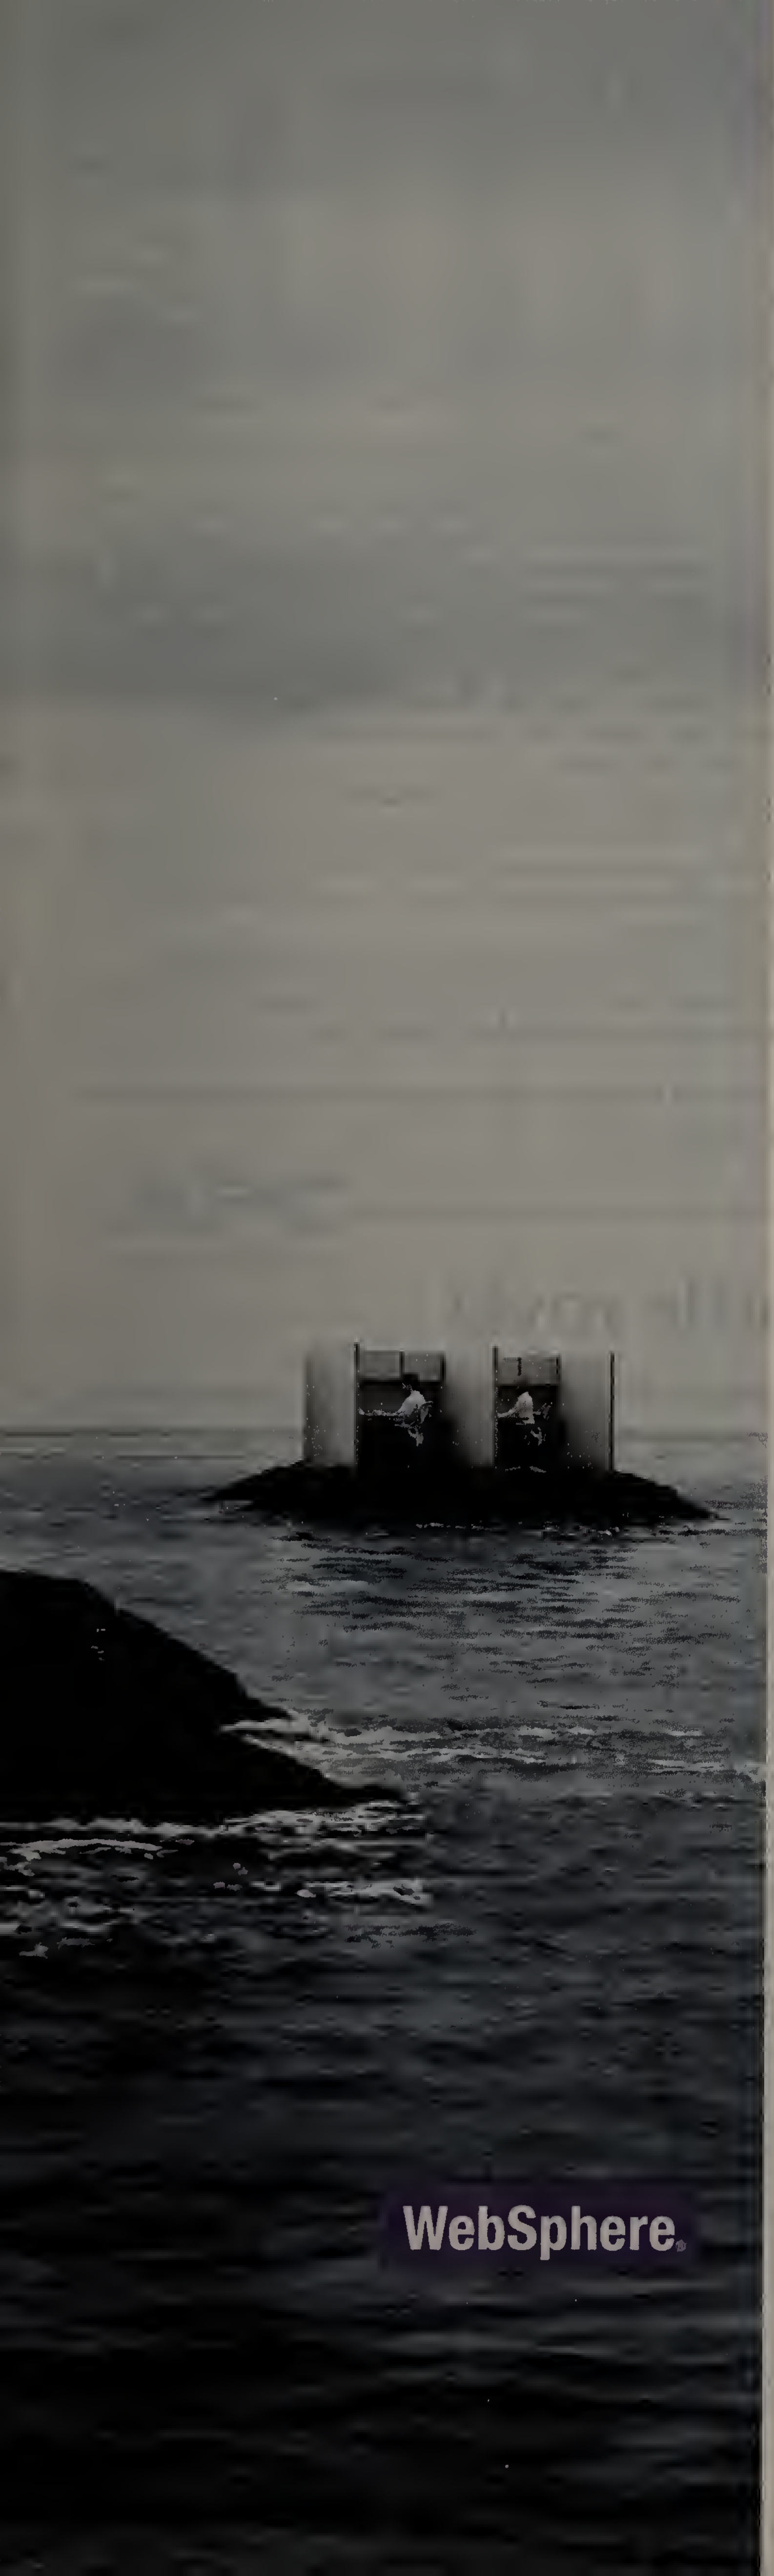
WebSphere.

Apple's iPhone — not the home of the free

NET INSIDER
Scott Bradner

Last week my mother admonished me for having published two columns about the Apple iPhone before it was released, but not a word since. She, of course, is right. I should have said something, but I've been trying to figure out what bothers me so much about the product.

I have not bought an iPhone — I may, but I'm not sure if or when. I have played with them and am astonished at their quality and ease of use. I expected a lot from the Apple designers, but until I held an iPhone and played with it, I had not internalized just how good a consumer product could be. The iPod should have given me a big hint (www.nwdocfinder.com/1425).

Apple also has surprised most of its competitors in the advanced phone business. A few are trying to put out iPhone clones, and a few of these devices look good, but I expect it will be a long time before products appear that show that other vendors understand anything about what Apple has done. Making a clone does not require understanding; you only have to look at the iPod to see how hard it has been for most vendors to "get it." Apple introduced the iPod in 2001 (www.nwdocfinder.com/1426), and to me, even today, there are no other products that come close to it in user-interface design. (And there are rumors that we may be just weeks away from a whole new iPod design, maybe something like a phoneless iPhone.)

So, the product itself is — as far as I can tell without living with one — great. According to the surveys I've read, most of the people who actually bought iPhones are very happy with them. The network managers in their companies may not be as happy because the iPhone is missing some things that such network managers see as required for an enterprise phone, including high-quality interaction with Microsoft e-mail systems and remote-device lock and erase.

There is a lot that bothers me about the iPhone, however, mostly about Apple's business decisions. Back in January, I wrote about some of the technology I'd like to see in the iPhone (www.nwdocfinder.com/1427). Most of what I wanted is not there. Lots of other things are, but the functions that would make the device complete are missing, at least from Apple. Some of the missing parts already are available from third parties. It is hard to blame Apple for not being able to lock out the hackers, especially when they have your device in their hands (www.nwdocfinder.com/1428), but to me, it would have been far better for Apple to sell a version of the iPhone that admits it is a computer running a good operating system and lets customers use it openly.

The worst part of the iPhone is that Apple is treating the iPhone just like another cell phone. Apple, the company whose innovative and compelling business model forced the music business and some of the TV and movie business to deal with the Internet, has done none of this when it comes to the iPhone. The phone, as sold in the United States, is locked into a particular carrier.

The locks, predictably, were quickly overcome and now Apple is retaliating by trying to block the exploits. If it were true to its image, Apple would have sold unlocked phones to people who wanted them. It may have to in Europe. If so, it will be sad indeed if customers in Apple's own country can't be free.

Disclaimer: Harvard predates the "land of the free" but has not expressed an opinion about Apple's refusal to be part of it in this case. Thus, the above review and lament are mine alone.

Bradner is Harvard University's technology security officer. He can be reached at sob@sobco.com.

Unified communications — battle royal

TOLLY ON TECHNOLOGY
Kevin Tolly

For me the recent VoiceCon show in San Francisco gave new meaning to the words "unified messaging." As I made my rounds to close to two dozen analyst meetings, almost every executive was focused on laying out his company's "Unified Communications" strategy and/or its upper-stack cousin, "Communications-enabled Business Processes." UC and CEBP were certainly the stars of the show but how we'll get there is not at all clear, and a big battle is brewing.

IP telephony battles of recent years have, naturally, been between the long-standing PBX vendors — Alcatel, Avaya, Nortel, Siemens — and the "new" VoIP vendors — Cisco, Shoretel and a string of others. Now, with its Office Communications Server 2007, Microsoft is arriving in a big way and has big plans to take over IP telephony, er, sorry, unified communications.

In an hour-long product commercial, apparently mislabeled as a keynote speech in the event program, a Microsoft executive spoke in detail of how the aforementioned OCS 2007 and the client-counterpart Office Communicator 2007 would, essentially, eliminate the need for old, "hardware-based" systems. Meaning, in essence, anything sold by anyone other than Microsoft.

As I sat there, I couldn't help but remind myself that there is really no such thing as a "hardware-based" PBX anymore. Years ago, PBX systems did run proprietary software usually on proprietary hardware but those days are over. The "traditional" vendors have all ported the most important system elements to run on open hardware and OS platforms. All offer "softphones" that run on popular PC clients, and VoIP "hardware" phones are more software than hardware.

As it happened, my meeting after this session was with executives from Avaya. When I asked them if Microsoft's PBX-elimination strategy concerned them, they said it did not. Why? Because they believe that CEBP will trump unified communications. To explain: Where unified communications provides integration of "general" functions — like being able to call someone by clicking on his name in Outlook — CEBP will provide a more significant value add.

CEBP, Avaya and others say, will allow communications tools to be directly integrated into business processes. Avaya's example (and another keynote topic) was Black and Decker. CEBP (from Avaya) lets the company use text-to-speech to allow computers instead of people to call customers whose products have been repaired and are ready for pickup.

While it is hard to argue with any of the individual approaches, they all can't win. In the past, the battlefield was usually limited. For example, when IP telephony came about the struggle was between old-line PBX "telephone" departments and the IT department. With server virtualization, the struggle is often between the server people and the data center infrastructure teams.

When it comes to unified communications/CEBP, the battlefield can run all the way from the IP telephony team, across to those responsible for the company messaging and server strategy all the way up to the application teams that program the company's line-of-business applications.

If the CEBP strategy can truly have the application drive the communications technology, then the traditional telephony vendors have a good shot. All too often, though, the IT infrastructure is put in place for subsequent use by application teams. If that happens this time, it could be an all Microsoft communications world before anyone notices.

Tolly is president and CEO of The Tolly Group. He can be reached at ktolly@tolly.com.

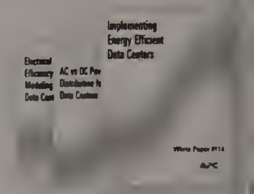
Implementing
Energy Efficient
Data Centers
Electrical
Efficiency
Modeling
Data Cent
AC vs DC Pow
Distribution fo
Data Centers

UTM firewalls: *READY FOR THE ENTERPRISE*

OUR TESTING SHOWS THAT UNIFIED THREAT MANAGEMENT APPLIANCES AREN'T JUST FOR THE SMB MARKET ANYMORE

BY JOEL SNYDER

IT managers at small and midsize businesses like unified threat management appliances — firewalls that layer on antimalware protection, content filtering, antispam and intrusion prevention — because deploying a single, multifunction device reduces costs and simplifies configuration.

However, deciding whether and where to deploy UTM appliances in a large enterprise is a more complicated and difficult decision. The idea of a single point through which all traffic flows as an obvious locus for threat mitigation doesn't work when a network has dozens, hundreds or thousands of distinct locations. Also, because performance is a critical issue in large networks, savvy network managers often seek to distribute threat protection rather than centralize it, simply to reduce the likelihood of a performance bottleneck.

Similarly, the style and quality of threat mitigation features one commonly sees in an SMB UTM may not be of interest to an enterprise, where requirements are more exacting and security architectures are more complex. For example, the antispam features and functionality in UTM firewalls pale compared with those in stand-alone enterprise-class dedicated antispam /antivirus appliances.

With such dramatic differences between SMB and enterprise requirements, is there a place for enterprise UTM firewalls? The answer is definitely "yes," for these three reasons: reduced complexity, simplified management and increased flexibility.

Reduced complexity

Enterprise network managers have long sought to include additional threat protection, especially intrusion detection/prevention systems (IDS/IPS) functions, both at the core and at the perimeters of their networks. However, the complexity of dropping stand-alone IDS/IPS boxes into a network has made them wary.

Building the "firewall sandwich," with load balancers surrounding a core of clustered firewalls, is well understood, but trying to scale that sandwich up with another layer of protection dramatically increases architectural complexity and potential instability.

A simple sandwich is considered science by network architects, but adding layers takes it from craft to art, dramatically increasing the difficulty of the project and opening a window for failure and problems. It's like adding just one more piece of cheese to that Dagwood sandwich: Not only will you be unable to get it in your mouth, but the whole thing may fall apart on your plate.

Enterprise UTM with integrated IDS/IPS can give network managers additional security throughout the network without the massive increase of complexity that stand-alone IPS devices would create.

It's pleasant to imagine the concept of a single UTM console that can handle everything from IP routing to IDS alerts, but enterprise security teams often want different management systems for a reason: different people are responsible for different kinds of threats and configuration.

Nevertheless, some level of management integration can reduce the task of handling these different functions. For example, every management console must have different network objects in it that are used to define policy: here are my mail servers, here are my users, this is the guest net-

GIACOMO MARCHESI

work, here is where the Internet is.

Each time those same objects must be typed into a different management system, and each time these objects are updated and adjusted, there is an opportunity for human error or miscommunication to create a security hole. A single management console that shares objects across different functions simplifies the complex task of management.

This single management view is especially valuable when firewall, VPN and IDS /IPS are considered together because all three of these functions act on the same policy. Each of these functions needs to have some view of the topology of the network, what applications are running on different servers and what different groups of users are allowed to do. Completely separate management for all three functions makes coordinated policy maintenance difficult, if not impossible.

A single UTM-ready management console realistically enables a fine-tuning of policy across all three functions, increasing total security.

Enterprise security architects generally scoff at the plethora of features, such as antivirus, antispam, antimalware and antiphishing, that are being built into SMB UTM devices. With a "best of breed" mentality and correspondingly large budgets, they are barely interested in activating IPS features in their existing firewalls. However, there are always specific situations where the ability to turn on, for example, antivirus, may be a huge benefit.

Having additional security features latent in large firewalls that can be activated with the click of a mouse gives the network manager increased flexibility, which is of significant value. For example, blocking incoming viruses in a UTM firewall may be a life-saver when the normal antivirus appliances suddenly stop working because of hardware, software or update failure.

Or consider the requirements of a guest user network: Most enterprises have chosen HTTP proxies to provide content filtering and antiphishing protection but may want to let guest users choose a different kind of protection and not take on the support burden of making sure they're properly working with the enterprise proxy. It may be simpler and more effective to enable these features in a UTM firewall for those networks.

The flexibility to bring security services in and out of the equation quickly using a UTM firewall supports threat response requirements — even if those features are rarely used.

Snyder is a senior partner at Opus One, a consulting firm in Tucson, Ariz. He can be reached at Joel.Snyder@opus1.com.

Enterprise UTM pros and cons

Pros:	Cons:
Complexity: High availability and scalability are dramatically simplified in UTM.	**Performance:** Enabling threat response features causes a huge performance hit and makes performance unpredictable.
Management: A single management interface enables better coverage for less effort, and reduces the possibility of mistakes.	**Choice:** Bundled threat response represents choices the vendor made based on partnerships and commercial interests, not necessarily matching what you'd choose for your own network.
Flexibility: Ability to bring security services in and out of the equation quickly supports threat response requirements best.	**Features:** Threat mitigation bundled into firewalls usually doesn't match the functionality and features in stand-alone products.
Cost: Long-term costs for UTM will likely be lower than individual point solutions.	**Separation:** Different teams are responsible for different threats, and requiring coordination and agreement between them can be difficult and time-consuming.

Top trends in enterprise UTM market

BY JOEL SNYDER

1. All firewalls are for unified threat management. There is little distinction between a UTM firewall and a "normal" firewall nowadays. The firewall vendor community has made the transition so that all current products include the option to include some UTM features. While very high-end devices may not include much beyond embedded intrusion-prevention systems and VPN, the term "UTM firewall" has become redundant. If it's a modern-day firewall, it does more than simply block or allow traffic.

2. Conversely, UTM doesn't necessarily include the firewall. Whether it's a public relations ploy or a search for more customers, the UTM market has expanded to include products that don't actually have a firewall inside. Several vendors have brought products to market that have weak or nonexistent firewalls, yet a strong suite of threat mitigation features, including antivirus, antimalware, content filtering and traffic analysis. By combining these everything-but-the-firewall features into a single system, such vendors are focusing on the threat mitigation features and can design hardware that fits those requirements best to bring a very strong offering to the table.

3. New products have new architectures. Most UTM firewalls do a poor job at certain functions — antispam and antivirus are the best examples — because the underlying hardware and software was not originally designed to meet the needs of UTM. For example, without disk space, a UTM firewall can't provide a spam and virus quarantine. Or, without a link to the corporate directory, user personalization and differentiation on settings can't occur. While established vendors are not moving quickly in this area, new products are coming to market that reflect a rethinking of software and hardware requirements for a UTM firewall that provide better coverage on the threat mitigation side of the house.

4. Vendor business models are evolving. UTM changes the model from a capital-focused one to a service-focused one. This means that firewalls will get even less expensive — but only be really useful when under a support agreement that provides constant updates. In fact, small-to-midsize-business-sized software-based firewalls are coming to market for "free," based on the idea that they will generate revenue through support and subscription fees. It worked for razors; it can work for firewalls.

5. Network managers remain skeptical. While vendors are packing features into products and offering them at attractive prices, network managers are still hesitating to enable threat mitigation features. The UTM sweet spot is networks in SMBs with no dedicated security staff. While you'd think that enabling UTM features is a no-brainer on these new devices, fears of false positives and bad experiences with performance slow-downs keep many of these devices running in firewall-only mode. Enterprise network managers are even further behind than their small-business brethren in deploying UTM features such as IPS in high-end devices.

How to select enterprise UTM firewalls

BY JOEL SNYDER, NETWORK WORLD LAB ALLIANCE

Selecting UTM firewalls in an enterprise environment is more work than just picking a standard firewall, because the "UTM" moniker doesn't offer much information about what the firewall actually does. When evaluating enterprise UTM firewalls, there are four key issues to consider: performance, UTM feature set, network integration and management. Many of these overlap traditional firewall requirements but must be considered in the light of specific needs for very high-reliability, high-performance, enterprise-class products.

Performance is the key starting point for UTM firewalls, because the UTM features exact such a heavy performance cost. Without accepted metrics on how to measure UTM firewall performance, network managers are left to determine how fast a UTM device will go by turning it on and putting it in the middle of their network. No matter what you do, don't skip this step or some reasonable approximation in a test lab. The performance of UTM devices is very dependent on exact configuration and traffic flows, and without some testing, you could easily end up with a device that can't handle the loads you throw at it.

UTM firewalls that let you scale up without a forklift upgrade, either by upgrading in the chassis or by adding systems in an active/active load-balancing configuration, are especially attractive when the performance card is on the table. But it's better to start with a system that can run as fast as you need the day you turn it on, and save upgrading for another year.

UTM features are near the top of the list for selection criteria. The idea seems simple enough: If you want antivirus, it should do antivirus. But within UTM firewalls, there's considerable variation in how a simple feature such as antivirus is implemented. For example, not every firewall can examine every protocol for virus signatures, and even those that do cover the top protocols can't

See UTM, page 40

FIVE TIPS ON DEPLOYING ENTERPRISE UTM

Early rounds of testing in our upcoming 10-vendor shootout of enterprise unified-threat-management firewalls have shown that deploying enterprise UTM has its own pitfalls. Here are some tips to help you avoid those issues in your network.

1. Don't try to do it all in one box.

Although you can buy UTM firewalls of almost unlimited power, that doesn't mean you should try and consolidate all your firewalls into a single system. It's important to logically distribute firewall functionality, because of the difficulty of building a single, coordinated, enterprisewide policy. Even though firewall vendors have made huge strides in centralized management, no product easily handles many zones of control with differing firewall rules, network address translation rules and VPN tunnels in a single policy. Add in the axes of intrusion detection/prevention systems (IDS/IPS) or other UTM features and the policy becomes even less manageable. UTM devices can support consolidation, but it's easy to go too far. Make sure you don't "over-consolidate" into an unmanageable device.

2. Check performance carefully.

Performance is one of the biggest gotchas in UTM devices: As you turn on features, performance can drop dramatically — or not at all. Security product vendors don't hide these performance costs, but they don't make it easy for you to understand what the impact of enabling different UTM features will be on your system performance. Make sure you know exactly what your UTM configuration will be, and test it to be sure that performance matches your requirements. Speed drops of 75% to 90% are common with a single check box. Be sure you also have plenty of headroom. IPS rules, for example, will only get more complex over time, so your IPS will get slower and slower over time.

3. Don't shortchange management.

UTM firewalls have a lot to say, with each layer of the firewall logging information about the traffic flowing through it. Enterprises are increasingly being asked to capture and retain these voluminous firewall logs for months or years. Make sure you plan for a dedicated management server with plenty of disk space, memory and CPU power to handle these chatty boxes. Although some enterprise vendors still allow management to be handled via a Web GUI or through a management server running co-resident with a firewall, don't be tempted to skip a properly separated and sized management system.

4. Verify high-availability and scalability features.

As firewalls take on more functions and become more central to correct network operation, ensuring high availability and scalability also is more important. Because performance is more likely to be a bottleneck in UTM, active/active configurations are more attractive than active/passive — but such configurations are more difficult to build and test. Simulating all the different failures, and making sure that you test them in all the different states of the cluster, can be a five-day and not a five-minute job. We also found that not every feature in our UTM devices works in the same way. For example, the basic firewall and VPN functions are usually shared cleanly across a cluster, but dynamic routing may not be as well thought out. If the VPN tunnels stay up across an individual device failure but the cluster doesn't know how to route the packets, that's not "highly available."

5. Complex configurations are hard to verify.

During our testing, we found that the firewalls often were not doing what we thought we had asked for, especially in the area of UTM add-ons such as antivirus and IPS. You should be prepared for a second round of training on system management and configuration, because what you thought you knew about your enterprise firewall may not be enough to get a proper UTM configuration in place. Even if you think you know what you're doing, it's valuable to run simple tests to validate that the protections you've asked for are actually activated. The terminology and protocol coverage varies wildly across different products, and a simple check box for a UTM feature may need an hour of testing to understand.

— Joel Snyder

DATA CENTERS

Resources to help users better support data centers

DATA CENTER CHALLENGES (% indicating "Very Challenging" or "Challenging")

53%	Troubleshooting software problems
50%	Maintaining disparate applications
48%	Issuing software patches
48%	Ensuring adequate performance and availability
41%	Safeguarding the data center from physical disaster
40%	Scaling the environment up and down for demand peaks and valleys
39%	Having enough physical space in the data center
34%	Adequately cooling equipment
30%	Understanding the interdependence of data center equipment
28%	Dealing with power outages
23%	Troubleshooting hardware problems
23%	Keeping track of the equipment in the data center

Today's challenges of supporting a data center include virtual server sprawl, ongoing migration to blade servers, mounting cooling demands, a never ending need for more power, the rising costs of energy and more.

Network World can help you alleviate these challenges with a collection of resources that offer concrete suggestions and plans of action.

Go to:
www.networkworld.com/DataCenterResearch
for all data center research.

NETWORKWORLD®

Five Strategies for Cutting Data Center Energy Costs Through Enhanced Cooling Efficiency

See how to optimize your data center efficiency through virtualization, digital system controls and emerging monitoring capabilities.

EMERSON
Network Power

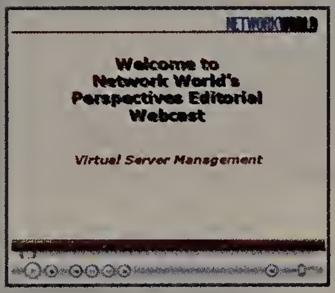

NETWORKWORLD

Welcome to Network World's Perspectives Editorial Webcast

Virtual Server Management

Network World Editorial Webcast: Virtual Server Management – Weighing the Options

Virtual server sprawl is a byproduct of virtualization. Discover new tools designed to help alleviate the management issues involved.

Gateway.

A Unified Approach to Workload Lifecycle Management

Find out why your organization should consider adopting a unified approach to managing workloads in the data center.

PLATESPIN

Best Practices to Control Your Data Center

Read about solutions that help IT shops better support remote data center maintenance with this in-depth whitepaper.

Avocent
The Power of Being There.

UTM
continued from page 38

always be configured to work on nonstandard ports. One firewall we tested only looks for viruses in certain defined Multi-purpose Internet Mail Extensions types as a way to keep performance peak, opening the potential for future exploits to slip directly past. A critical exercise before buying is understanding exactly what coverage is included and how that coverage relates to your own traffic patterns and requirements.

A small number of UTM firewalls offer a choice in threat mitigation products, such as multiple antivirus vendors, but most lock you into a single vendor. While antivirus (as an example) is considered a commodity service, other services, such as IPS and antimalware, are in more active development — which makes the choice of vendor and consistency of implementation significantly more important.

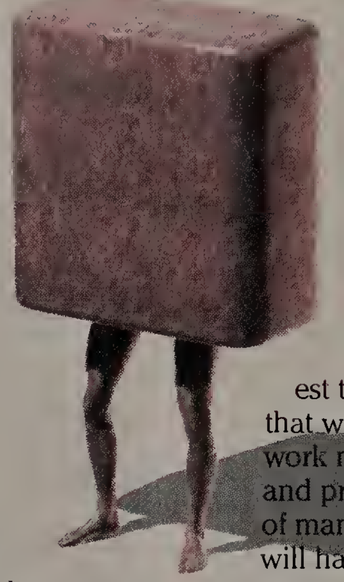

Network integration includes the aspects of a UTM firewall that let it sit securely within an existing network. For example, enterprise UTM firewalls are more likely to need some support for dynamic routing protocols such as Open Shortest Path First to integrate within an existing infrastructure. Virtual LAN support, high port density, WAN support and expandability of interfaces over time are all similar network integration features. While most of these also are relevant in a pure enterprise firewall without UTM features, the tendency to use UTM firewalls as points of consolidation of security control raises their importance.

Another aspect of network integration includes the equipment and interfaces required for high availability and scalability. If you've got a specific set of load balancers or switches, the UTM firewalls have to be able to integrate with that equipment with a minimum of reengineering and additional equipment. Similarly, with the additional requirements for active/active clustering that UTM performance brings, full support for upward scalability should be considered a UTM evaluation criterion.

Management is one of the most difficult parts of a UTM firewall to evaluate, because you don't know how good or bad the management is until you've had lots of experience with the product. While most management systems strive for glitzy interfaces for the novice, the real evaluation comes with consistent and continued use. Unfortunately, by that time, it's too late to choose another product.

In UTM products, one of the most important features of management is the ability to bring UTM features into play in a flexible and controlled way. For example, a management system that requires all traffic to flow through the IPS, or none of it, is not suitable for an enterprise UTM device. At the same time, the management system must allow for different profiles for the same UTM feature. For example, an IPS might be configured in a liberal way for internal users browsing the Internet, while turned up to strict levels for guest users coming from a different subnet.

While UTM management systems will be mostly of interest to the security manager, there are aspects of configuration that will fall to a desktop manager (such as antivirus) or network manager (such as dynamic routing). Separating function and privilege level horizontally and vertically across the domain of management is difficult. However, if your UTM deployment will have people from three (or more) teams peering into the same management system, features in this area can be critical to successful long-term operation. ■

ONLINE: UTM Buyer's Guide

This Buyer's Guide comprises intrusion-detection systems — offered either as stand-alone products or as components of Unified Threat Management devices — that detect possible network breaches using either signature- or anomaly-based techniques. **www.nwdocfinder.com/1069**

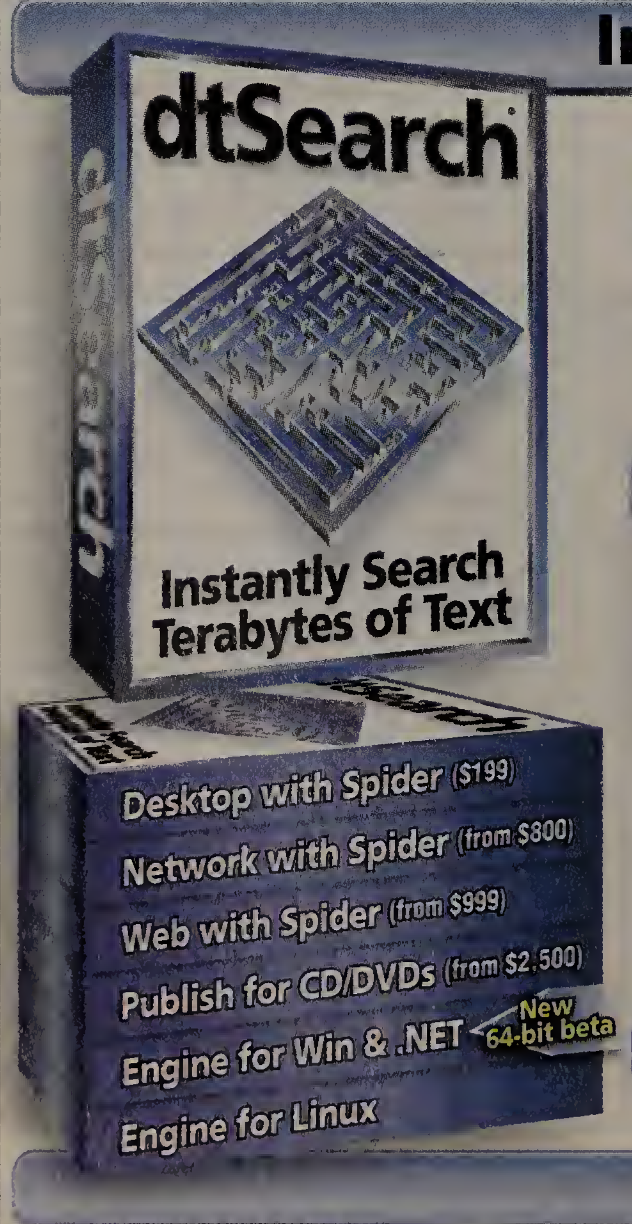
dtSearch
Instantly Search Terabytes of Text
Desktop with Spider ($199)
Network with Spider (from $800)
Web with Spider (from $999)
Publish for CD/DVDs (from $2,500)
Engine for Win & .NET
New 64-bit beta
Engine for Linux

Journey to the storage frontier.

Storage Networking World. Dallas.

October 15-18, 2007
Gaylord Texan Resort
Dallas, Texas

Storage Networking World—the world's largest and foremost storage networking event—is where IT management and professionals learn, network and maximize their company's storage capabilities. At SNW, you can choose from over 140 educational sessions and network with peers from around the globe—plus visit with top solutions providers in the world's largest Expo, Solutions Center and Hands-on Lab focused on storage. SNW's vast agenda makes it the most appealing and well-respected assembly of storage networking professionals in the world. This is your opportunity to amass reliable, firsthand, practical knowledge in only a few days, set against a beautiful Texas backdrop.

See Frank Abagnale

The world's foremost authority on fraud and identity theft, and the inspiration behind the now-classic book and film, "Catch Me If You Can," Frank Abagnale, is the "con man who came in from the cold," agreeing to help the FBI as an expert on fraud, after a five-year prison term. Abagnale's lecture was hailed by Tom Hanks as "the best one-man show you will ever see."

See V. James Onalfo

As CIO and Deputy Commissioner of the NYPD, Jim has received six awards for the Real Time Crime Center, including the Top Government Global award. Jim accepted the Computerworld Honors Program's 21st Century Achievement Award in 2006, and was recently recognized as one of Computerworld's Premier 100 IT Leaders for 2007.

PARTNER PAVILION SPONSOR

EMC² where information lives®
LSI
Microsoft®

PLATINUM SPONSORS

BROCADE
CISCO.
datadomain
DELL™
EMC² where information lives®
HITACHI Inspire the Next
hp invent
IBM
intel
LSI
Microsoft®
NetApp®
PMC PMC-SIERRA
QLOGIC®
Seagate
Xiotech®

GOLD SPONSORS

3PAR
adaptec
AMD
BLUEARC®
ciena
commvault
compellent
COPAN SYSTEMS
CROSSROADS®
dot HILL®
:::fast
FCIA
Finisar
H3C IToIP Solutions Expert
IBRIX®
intransa®
IRON MOUNTAIN®
ISILON SYSTEMS
PROMISE TECHNOLOGY, INC.
Quantum
SPECTRA
SUPERMICRO®
Verari systems®
xyratex

CONTRIBUTING SPONSORS

acopia NETWORKS
agámi systems
alacritech
APTARE
arkeia SOFTWARE
attune
CipherMax
DataDirect NETWORKS Performance. Capacity. Innovation.
EQUALLOGIC
FalconStor Software
HIFN
INDEX ENGINES
Infortrend
InMage
Z-Star Networks, Inc.
NEC
NETEX
NETGEAR®
NEWISYS®
ONStor
OPSWARE INC
PACKETEER®
panasas
PERMABIT
Plasmon
PROSTOR SYSTEMS
RELDATA
riverbed
SEPATON®
ServerEngines™
sgi
Storewiz
TMS TEXAS MEMORY SYSTEMS
WysDM Software

SNIA

STORAGE NETWORKING WORLD
COMPUTERWORLD

October 15-18
Dallas, Texas

For complete details, or to register: www.snwusa.com

IBM Lotus Sametime serves up messaging any way you want it

Jabber and Cisco follow as close seconds in test of corporate IM platforms

BY BARRY NANCE, NETWORK WORLD LAB ALLIANCE

Messaging has come a long way from the early days of rudimentary chat programs, the DOS and Windows "NET SEND" command and the Novell NetWare "SEND" command.

The ideal corporate instant-messaging environment lets users communicate anything they choose, from simple typed messages to documents to video. It tells employees which colleagues are available for an impromptu meeting and which don't wish to be disturbed. The ideal IM environment offers impenetrable security that thwarts intrusion attempts, as well as IM-borne malware. It's nimble and responsive; intuitive to use and administer; and integrates seamlessly with other IM products and protocols, such as AOL Instant Messenger (AIM).

Preferably, it safely archives IM sessions for easy retrieval by an auditor, is highly scalable, exhibits rock-solid reliability and uses network resources frugally. A corporate IM product taps into a Windows Active Directory or a Lightweight Directory Access Protocol (LDAP) back end for grouping and authenticating users. And finally, it provides the necessary VoIP capabilities to turn a chat session easily into a telephone call.

In short, the model platform makes holding meetings via IM as productive as — or even better than — meeting face to face.

To test the state of corporate IM tools we invited all vendors in this space to send products. We received Extensible Communications Platform (XCP) 5.2 from Jabber, Lotus Sametime 7.5.1 from IBM and Openfire Enterprise Edition 3.2 from Jive Software. We downloaded Gordano Messaging Suite (GMS) 5.0 from Gordano's FTP site and Mirador Instant Messenger for Windows 3.0 from Serial Scientific International's (SSI) Web site, and we accessed Cisco's WebEx AIM Pro Business Edition via the Internet (see "How we did it," page 44).

IBM Lotus Sametime earned our Clear Choice Award for its superior messaging, high level of integration with other applications, ease of use, scalability and excellent security. Nearly as excellent and carrying a much lower price tag is Jabber's XCP. Cisco's WebEx AIM Pro is a great choice if you prefer to outsource server operations and your users have reliable Internet connections.

IBM Lotus Sametime

Sametime is a feature-rich environment for network-based collaboration and conferencing. It consists of the Sametime Server and client-based Sametime Connect software. Users can message each other via Sametime Connect or a Web browser, or from within Lotus Notes. Sametime Connect also can be launched directly and easily from within Microsoft Office and Outlook. All these points of entry worked well in the lab.

Sametime's messaging interoperated seamlessly via IBM-supplied gateways with AIM, GoogleTalk and XCP. Setting up these gateways involved installing the software on Internet-accessible servers and, in the case of AIM, installing a digital certificate to authorize the IM traffic.

Sametime's security used 128-bit encryption for data privacy, and users were authenticated against LDAP, or Lotus Domino servers if we specified. Our Sametime hacking attacks — which included robot password crackers and, for eavesdropping, protocol analyzers — were futile. Sametime also kept IM-borne spyware and spam from annoying our users. Furthermore, IBM says it soon will change its encryption method to be compliant with the Federal Information Processing Standard 140.

Its reliance on LDAP or Domino for user authentication made administering Sametime simple. For example, we only had to publish the Sametime server's name, set up policies to allow or disallow file transfers, specify which users couldn't use the AOL gateway, specify the number of days to save IM transcripts and set a maximum image size for IM-transmitted screen captures. Additionally, Sametime let us search the IM archive by date or user for auditing purposes.

In our stress tests, Sametime never used more than 8% of the available bandwidth, which made it nearly as resource frugal as Jabber's XCP. IBM uses Sametime internally and claims it needs only four servers to support its 380,000 worldwide employees, who send 5 million messages each day.

Sametime's set of features is rich yet child's play to use, mostly because

NETRESULTS

Product	Lotus Sametime 7.5.1	Extensible Communications Platform (XCP) 5.2	WebEx AIM Pro Business Edition
Vendor	IBM www.ibm.com/lotus/sametime	Jabber www.jabber.com	Cisco www.webex.com
Price	$56.75 per user.	$35 per user.	$5 per user, per month (subscription).
Pros	Feature-rich, secure, well integrated with other applications, such as Microsoft Office.	Interoperable with AOL Instant Messenger, secure, scalable.	Internet-based service (if you prefer outsourcing), good Active Directory integration, good security.
Cons	Pricey	Needs to integrate better with Outlook and other mail readers.	Probably not for you if you have slow or unreliable Internet links, or you dislike outsourcing.
Score	4.9	4.6	4.5

SCORECARD

Action	Weight	IBM Lotus Sametime	Jabber Extensible Communications Platform	Cisco WebEx AIM Pro Business Edition	Jive Software Openfire Enterprise Edition	Gordano Messaging Suite	Serial Scientific Mirador Instant Messenger for Windows
Messaging	20%	5	5	4	4	4	3
Security	20%	5	5	5	3	4	3
Ease of use	20%	5	4	4	3	3	3
Interoperability	20%	5	5	5	4	2	3
Special features	10%	5	4	4	3	3	3
Documentation/installation	10%	4	4	5	3	3	2
Total score		4.9	4.6	4.5	3.4	3.2	2.9

Scoring key: 5: Exceptional; **4:** Very good; **3:** Average; **2:** Below average; **1:** Subpar or not available.

IBM gave it user-oriented conveniences. For example, Sametime changes a user's "presence" automatically to "in a meeting" when the user's Notes or Outlook calendar indicates there is a meeting scheduled. When users are away from their PCs for a specifiable period of time, Sametime automatically marks their presence as "away." And it adds a system tray icon that makes changing presence quick and painless. Sametime's presence concept, in addition to denoting that a user is away, busy or in a meeting, reveals geographic-location data, so users know colleagues' time zones. It even lets users specify they are available to some users but busy to others. Going beyond text messaging to share documents, images and video is easy in Sametime, and it integrates with VoIP to make switching from typed messages to a phone conversation (multiple party, if you like) completely transparent.

Sametime's Web conferencing automatically captures details of who attended a meeting and a transcript of the meeting. It offers breakout sessions within the overall Web conference, and users can tell Sametime to switch to off-the-record mode to prevent anyone from saving information they've typed but don't want attributed to them.

The Sametime server software, which requires that Lotus Domino be installed, runs on IBM's AIX and i5/OS, Linux (Red Hat and Novell's SUSE), Sun Solaris, and Microsoft Windows Server 2000 and 2003.

Extending Sametime with custom programming to integrate, for example, with an in-house written application is easy through its well-documented programming interface. With less than a day's programming, we added Sametime awareness via presence and contact names to a Visual Basic program.

Sametime's copious printed documentation is clear and comprehensive, and even includes a "Sametime for Dummies" booklet. Installation took less than an hour.

Jabber Extensible Communications Platform

XCP had an impressive range of features; scaled extremely well in a linear fashion; and integrated well with other IM environments, such as AIM (via Jabber's AIM Gateway) and Lotus Sametime (via a Sametime gateway).

XCP consists of a Connection Manager, Jabber Session Manager and Core Router. Client connections, gateways and server-to-server connections go through Connection Managers. The Jabber Session Manager processes sessions for individual clients, as well as presence and roster data. All components communicate through Core Routers.

The server software runs on Windows Server 2000 and 2003, Red Hat Linux and Solaris.

Jabber's platform authenticates users rigorously. XCP exhibits excellent security with respect to authentication and confidentiality. Using Simple Authentication and Security Layer technology, XCP verifies the identity of each client. Because the XCP server validates ("stamps") sender addresses, hackers can't spoof addresses to insert themselves into the XCP environment. And Transport Layer Security (TLS) ensures no eavesdrop-

Openfire Enterprise Edition 3.2	Gordano Messaging Suite (GMS) 5.0	Mirador Instant Messenger for Windows 3.0
Jive Software www.jivesoftware.com	Gordano www.gordano.com	Serial Scientific International www.e-securion.com
$15 per user.	GMS Instant Messaging, $450; GMS Collaboration, $950; GMS Mail, $450; and GMS Archive, $1,110. All prices listed are for 25 users.	Starts at $335 for 10 users.
Excellent "chat with an agent now" instant-messaging environment.	Presence includes geographic location, good integration with Outlook.	Excellent remote-control tool, switches easily between IM and VoIP conversations.
Not highly scalable.	Not interoperable with other IM environments (by design).	Windows-centric, documentation too brief.
3.4	**3.2**	**2.9**

ping of messages occurs. XCP even blocked spyware and IM spam.

XCP stores registration, authentication, user lists, electronic business cards and offline message data in an Oracle database (supplied by the IT department); and it can access user data stored in LDAP or Active Directory repositories. We tested the Oracle storage option, which was easy to set up and use.

XCP uses XML within the Extensible Messaging and Presence Protocol (XMPP) to send and receive messages. We were able to efficiently and easily exchange messages and files, including video, through XCP's IM environment. XCP's VoIP integration, which let us switch from keyboard to voice and back again, also worked well. Because GoogleTalk also is based on XMPP, XCP clients can send and receive messages to and from GoogleTalk clients without needing a separate gateway. In our tests, XCP communicated seamlessly with GoogleTalk and AIM via the included Session Initiation Protocol/SIP for Instant Messaging and Presence Leveraging Extensions (SIP/SIMPLE) gateway.

XCP's browser-based administrative console was intuitive to navigate and responsive. We used it to authorize users and groups for access to the AIM gateway, monitor the running of Connection Managers, and specify the severity level of XCP syslog entries. Simulating an audit, we searched XCP's message archive by date and user to examine the content of IM sessions.

Jabber claims that a single XCP server, configured with a pool of Connection Managers for controlling client/server sessions and linked to a single Oracle server, can support 2 million subscribers and 100,000 concurrent sessions with a latency of less than 0.29 seconds. Our stress tests, which subjected XCP to a barrage of messages from a simulated 1,000 clients, showed XCP used a meager 6% to 7% of available bandwidth.

XCP users can set their presence, which is displayed next to each contact name, to available, away or do not disturb.

Launching an XCP-based Web conference in the lab was a breeze. XCP interfaced easily with Adobe Acrobat Connect Professional, Cisco Unified MeetingPlace and WebEx. For mobile users, Jabber offers a client module for Research In Motion BlackBerry users, which also worked well.

XCP comes with a comprehensive programming interface for customers who want to customize or extend XCP's capabilities. The clear, easy-to-follow soup-to-nuts documentation is in printed form, and installation is a snap.

WebEx AIM Pro Business Edition

WebEx (purchased by Cisco last March) maintains IM servers to which corporate users can connect via a browser-based client module over the Internet. From anywhere on the Internet, you can log onto WebEx AIM Pro and chat with other employees or business partners.

WebEx handles all the messy details of server operation, such as monitoring utilization and making sure servers are up and running. While this can be a big advantage for customers who like to outsource server operations, it also can be a disadvantage. We had to trust WebEx to make its IM services always available and safely make backup copies of IM session archives.

Corporate Internet connections must be alive and well to use WebEx AIM Pro. To share files (especially video streams) reasonably fast Internet links (512Kbps or faster) are needed. Moreover, if some employees lack Internet connections — perhaps they're insulated from public access for security purposes — they won't be able to use WebEx AIM Pro.

WebEx AIM Pro works closely with WebEx's other offerings, such as the vendor's primary product, Web-based conferencing. Launching a WebEx conference session from within the IM client took just one mouse click. WebEx AIM Pro integrated with our Outlook calendars and address books to know, for example, when a person was in a meeting or otherwise away from his desk. From within a messaging session, we could share documents and even video clips easily. It also supports switching instantly from a messaging session to a VoIP-based phone conversation. Via WebEx-maintained gateway servers, WebEx AIM Pro gave us seamless access to AIM users.

We particularly appreciated WebEx's tools for batch uploading of user and group data from our Active Directory tree, and we could use our Outlook address books to initiate WebEx AIM Pro sessions as if the contacts were already in it. WebEx maintains message archives that administrators can search and download to ensure compliance with applicable laws.

Security consists of 128-bit SSL encryption for confidentiality, as well as password-challenge authentication by WebEx. The WebEx IM servers automatically scan messaging traffic for viruses, worms and other malware. They also block IM-based spam — unsolicited messaging sessions initiated from outside your network.

WebEx's online documentation is clear and comprehensive, and installation of the client module is a snap — no server installation is needed.

Jive Software Openfire Enterprise Edition

Openfire Enterprise Edition (formerly called CrossFire) is a commercial version of the open source Openfire server software. The Enterprise Edition — which requires Java 5 support and typically runs on Windows XP/2000/2003, Linux, Solaris and Apple's Mac OS X — adds such features to the open source version as a Web client, SIP softphone, more sophisticated reporting, better client management, message bookmarking and message archiving.

The commercial version also sports what Jive Software terms Spark Skinning, which lets users customize the look and feel of the chat client; and Fastpath, which automatically routes chat requests to the next available agent. Fastpath impressed us as we used it to transfer chat sessions, invite others to join a chat, set up canned responses and maintain a chat history.

Administering Openfire was painless. We viewed statistics on active users and conversations, monitored group chat rooms and searched through message archives by date, user and keywords. We created what Jive Software calls chat bookmarks, which tell users about each chat room's purpose and subject matter. We applied these bookmarks at our option to individual users, groups or all users. Openfire uses a published database schema and includes an embedded database. We used

See IM, page 46

How we did it

We evaluated each product's instant messaging capabilities, responsiveness, ease of use and ability to integrate with other IM products and protocols, such as AOL Instant Messenger. We wanted the products to integrate with Windows Active Directory and Lightweight Directory Access Protocol. Archiving messages for auditing purposes was a key criterion. To gauge security, we measured each product's ability to identify and thwart the sending of malware to IM clients, and we tested the products' ability to securely identify and authenticate users appropriately. We tested each product's special features, as well as its VoIP and presence capabilities. We also looked for scalability, reliability, low network-resource consumption, ease of installation and documentation quality.

Virtually all our testing took place across 512Kbps frame relay, T-1 and T-3 WAN links. The test bed network consisted of six Fast Ethernet subnet domains routed by Cisco routers. Our lab's 50 clients used computing platforms that included Windows NT/98/2000/2003/ME/XP/Vista, Red Hat Linux and Mac OS X. The relational databases on the network were Oracle 8i, IBM DB2 Universal Database, Sybase Adaptive Server 12.5 and Microsoft SQL Server 2000. The network also contained three Web servers (Microsoft Internet Information Server, Netscape Enterprise Server and Apache), three e-mail servers (Exchange, Notes and Sendmail) and two file servers (Windows 2003 Advanced Server and Novell NetWare).

A Compaq ProLiant ML570 computer with four 900MHz CPUs, 2GB RAM and 135GB hard disks, running Windows 2000 Advanced Server, Windows 2003 Advanced Server and at other times, Red Hat Enterprise Linux, was our test platform for all the products' server components. We tested Jabber's Extensible Communications Server 5.2 on Red Hat Linux (a vendor recommendation) and all the other products' server components on Windows Server.

IBM
PAINT

continued from page 44

the schema to connect Openfire to Oracle; Jive Software says you also can use MySQL, SQL Server, Postgres, DB2 or Sybase Adaptive Server.

Openfire is XMPP-based and interoperates easily with other IM environments, such as XCP and GoogleTalk. Openfire includes a public-gateway software module so users can have messaging sessions with AIM users, for example.

Openfire's Java underpinning limits its performance and scalability. Our stress tests revealed that although Openfire's network use was less than 10%, it consumed considerable server CPU — 40% to 70%.

Openfire's security relies on the provisions within XMPP (primarily TLS), and the Openfire server makes certificate management a simple affair. With a little programming and setup effort, we linked Openfire to an LDAP server and to Active Directory. Jive Software says Openfire also can use native Windows or Unix Pluggable Authentication Modules authentication.

Jive Software's presence flag, which appears in the Web client's contact list, tells you whether another person is online, offline or typing. The IM Web client is a snap to navigate and use, and the bookmarks make finding the right chat room a breeze.

We found Openfire best suited for the sort of Web-based customer interaction that uses links that say, "Click now to chat with an agent." For example, in one test, we used Openfire's Fastpath to route chat requests efficiently to a pool of agents waiting for customer queries. It's less useful for intracompany employee conferencing and collaboration. To its credit, however, Openfire integrated with Microsoft Outlook's calendar, and its VoIP integration let us turn a messaging session into a phone call with a single mouse click.

Its online documentation is comprehensive but lacking in detail with respect to some server operations. Installation takes just a few minutes.

Gordano Messaging Suite

You can pick and choose the IM features you want to deploy across your network from this suite of well-integrated software components. We tested GMS Instant Messaging (GMS IM), the cornerstone module, as well as GMS Collaboration, GMS Mail, GMS Anti-Spam and GMS Archive. These are optional modules that added conferencing, e-mail, avoidance of unacceptable topics and message storage to our IM environment. The suite runs on Windows NT/SP/2000, Solaris, AIX and Linux.

GMS IM offers a native Windows client and a Java-based client. Their look and feel are similar, and both worked well in the lab. With each client, we opened chat sessions, sent messages, managed our contact lists and worked on documents with other users via GMS Collaboration. Gordano's presence flag, which appears in either client's contact list, informs you whether a contact is online and — when used with Microsoft Outlook's calendar — whether the person is in a meeting. GMS IM also shows location information based on IP-address geolocation (knowing where on a network particular IP addresses are located). GMS IM lacks VoIP integration.

In addition to directly launching the Windows or Java-based IM client to begin an IM session, a user also can start the Windows IM client from within Outlook, or the Java IM client from within Gordano WebMail. Via GMS Archive, GMS IM stored transcripts of our test IM sessions and e-mailed the transcript at our request to all session participants at the session's conclusion.

GMS administration is rudimentary. For example, the GMS console did not show us real-time traffic statistics that we could use to monitor IM activity. And we had to write a custom program to search the archives to audit for IM content.

GMS IM used a moderate 9% to 12% of network bandwidth during our stress tests.

GMS IM's security consists of transaction (session) logging, which let us investigate IM hacking attempts by searching the logs for unauthorized users. GMS IM's native Windows authentication and Active Directory authentication worked well in the lab. GMS IM also incorporates a virus filter and a spam filter, both of which thwarted our attempts to attack it.

Gordano deliberately engineered GMS IM to not work with other IM environments, such as AIM and GoogleTalk. The company says this approach helps its corporate customers keep employees from chatting with friends and family while at work. However, unless the company sets firewall rules against it, employees can still access AIM or GoogleTalk as separate, nonauthorized applications.

The online documentation is unremarkable, and installation takes less than an hour.

Mirador Instant Messenger for Windows

Geared especially to Windows-centric companies, MIM consists of a server component that runs on Windows 2000/2003/XP Pro and a client component that runs on Windows 98/ME/2000/XP.

We used MIM's central console to set up IM users and passwords, group users by department, search the IM archive by date and user, and view current IM activity levels. The central console also let us configure clients by individual user or group to allow or disallow starting a remote control session or a document collaboration session. We also could set a maximum message size.

Using MIM for messaging is straightforward. A user clicks on another user's contact-list entry to initiate a chat, which MIM then establishes if the target's presence flag is set to available (other values are busy and offline). Once in a chat session, a user can start MIM's remote control feature or transfer files to other users, if these actions are authorized by the administrator. Besides contact-based messaging, MIM lets users switch from messaging to VoIP-based conversations, and has a feature the company terms co-browsing — distributing office documents or Web pages to other session participants and collaborating on changes to those documents. This worked well because Microsoft Office versions 2003 and later support online collaboration. MIM's remote control feature was especially useful in online training sessions.

MIM's network use was 8% in our stress tests.

MIM authenticates users against its own internally maintained user list. Its security features let us restrict the file types circulated, and MIM includes a message audit feature that helped reveal the contact names of people who attempted to compromise the IM environment. We also could limit the IP address ranges of MIM clients to ensure access only by users known to be on our network. For the sake of confidentiality, MIM uses SSL to encrypt messages. However, it lacked virus, spyware and messaging-spam filters.

MIM's online documentation is too brief to guide administrators and users through all the product's functions. Installation takes a few minutes.

Conclusion

We unreservedly and heartily recommend IBM Lotus Sametime for IM in a corporate setting. It is feature-rich, intuitive to use, highly scalable and platform neutral. Jabber's high-quality XCP also is worth investigating, especially because of its lower pricing. If you want to put a "chat with an agent now" link on your company's Web site, Jive Software's Openfire may be just what the doctor ordered. With WebEx AIM Pro, you can outsource IM server operation and still get a full-featured IM environment.

Nance runs Network Testing Labs and is the author of Introduction to Networking, 4th Edition, *and* Client/Server LAN Programming. *He can be reached at barryn@erols.com.*

■ **Barry Nance** is a member of the Network World Lab Alliance, a cooperative of the premier testers in the network industry, each bringing to bear years of practical experience on every test. For more Lab Alliance information, including what it takes to become a partner, go to **www.networkworld.com/alliance.**

Other members: Mandy Andress, ArcSec; John Bass, Centennial Networking; Travis Berkley, University of Kansas; Jeffrey Fritz, University of California, San Francisco; James Gaskin, Gaskin Computing Services; Thomas Henderson, ExtremeLabs; Miercom, network consultancy and product test center; Christine Perey, Perey Research & Consulting; David Newman, Network Test; Thomas Powell, PINT; Joel Snyder, Opus One; Rodney Thayer, Canola & Jones.

IBM®

Help make your data center green with IBM Cool Blue™ technologies and energy management services.

Go green with virtualization: *Advanced server and storage virtualization from IBM can help you decrease your number of boxes and lower your energy usage.*

Go green with energy management: *IBM Systems Director can give you active energy management to help you track and cap your power consumption.[1] It can help you see and regulate how much power the systems in your data center are really using.*

Go green with more efficient systems: *IBM POWER6™ processors with Advanced Power Virtualization mean your systems can use less energy doing the same amount of work.[2] For instance, consolidating 30 Sun V890s into one rack of POWER6-based IBM System p™ 570s can save you over $100K a year in energy costs alone.[3]*

Go green with IBM: *IBM Services can help design your data center, holistically, for better energy usage. With outstanding technology and people who understand what that technology can do for your business, IBM can help make your data center green.*

Learn how to make your data center more efficient.
IBM.COM/**TAKEBACKCONTROL**/GREEN

[1] Currently available on IBM System x and IBM BladeCenter servers. Expected to be available on IBM System i and System p servers 11/07. Energy management capabilities of IBM Systems Director are not available on IBM System z. [2] Advanced Power Virtualization is optional and available at an additional charge. [3] For complete details, go to ibm.com/takebackcontrol/claim. IBM, the IBM logo, Cool Blue, POWER6, System p, Take Back Control, System x, BladeCenter, System i and System z are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries. ©2007 IBM Corporation. All rights reserved.

Start with the right rack, and you can't go wrong.

Get the seamlessly integrated, fully compatible NetShelter® rack system from APC®.

APC, the name you trust for power protection, also offers a comprehensive line of non-proprietary racks, rack accessories and management tools that ensure the highest availability in a multi-vendor environment. With APC racks, accessories, and management tools, you can design a comprehensive rack solution that meets your availability needs for today and that easily scales up for tomorrow.

Need assistance? Our expert Configure-to-Order Team can custom tailor a complete rack-mount solution that suits your specific requirements.

Contact APC today and protect your rack application with Legendary Reliability®.

The NetShelter® SX is vendor neutral and carries the "Fits Like a Glove" compatibility guarantee.

HP/COMPAQ • SUN • IBM
GUARANTEED COMPATIBILITY
DELL • CISCO • LUCENT

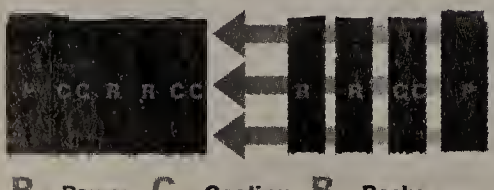

P = Power C = Cooling R = Racks

NetShelter is completely compatible with all APC award-winning InfraStruXure® architecture, allowing you to add rack, power and cooling on a scalable as-needed basis.

NetShelter® SX *starts at $1150*

Rack enclosures with advanced cooling, power distribution, and cable management for server and networking applications in IT environments.

- *Integrated rear cable management channels allow easy routing, management and access to large numbers of data cables.*
- *3000 lbs. weight capacity.*
- *Vendor neutral mounting for guaranteed compatibility.*
- *Toolless mounting increases speed of deployment.*

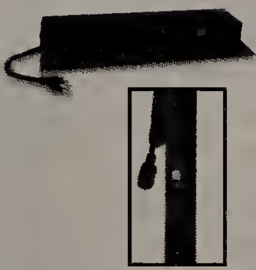

Rack PDU *starts at $89.99*

Power distribution that remotely controls power to individual outlets and monitors the aggregate power consumption.

- *Switched, metered, and basic models available.*
- *Includes horizontal, vertical, and toolless mount.*
- *Puts power in the racks near the equipment where it is needed most.*
- *Wide range of input and output connections from single-phase to 3-phase.*

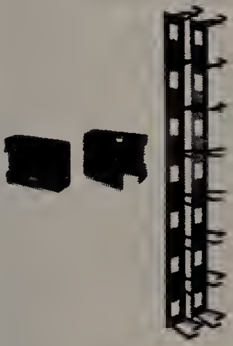

Cable Management *starts at $29.99*

Comprehensive selection of accessories designed to organize power or data cables within a rack environment.

- *Eliminates clutter and cable stress.*
- *Zero U of rack space with the vertical cable organizer.*
- *Quick-release tabs, toolless mounting.*

Rack-mount Keyboard Monitor *starts at $1550*

1U rack-mountable integrated keyboard, monitor and mouse.

- *15" or 17" ultra-thin, LCD monitor with integrated keyboard.*
- *Ease of installation minimizes support and maintenance costs ensuring lower cost of ownership.*
- *Can be used in a variety of IT environments from computer rooms to large data centers.*

Rack Air Removal Unit SX *starts at $2600*

Rear-door fan system for performance heat removal up to 23kW

- *Temperature controlled, variable speed fans allow reduced energy consumption during off-peak cooling periods.*
- *Ducted exhaust system increases air conditioning efficiency and prevents hot spots by eliminating recirculation.*
- *Manageable via Web, SNMP, Telnet and local LCD display.*

NetBotz® Security and Environmental *starts at $889*

Protecting IT assets from physical threats.

- *Visual monitoring of all activities in the data center or wiring closet.*
- *Third-party monitoring via dry-contacts, SNMP, IPMI, 0–5V and 4–20mA.*
- *User-configurable alarm and escalation policies.*
- *Temperature, humidity, and leak detection.*

Download Free Rack White Papers

For full details, Visit www.apc.com/promo Key Code x242x

• Call 888.289.APCC x9162 • Fax 401.788.2797

APC
Legendary Reliability®

132 Fairgrounds Road, West Kingston, RI 02892 USA AX4A6BF_NAMf

Tools for Remote Network Management

Access your Network Equipment from Anywhere!

When it comes to Remote Network Management, no one has more remote access tools than Western Telematic! Our products offer the design flexibility you need to mix and match equipment for small or large scale remote management strategies. WTI products are installed in thousands of network sites worldwide.

- ✔ *In-House Design & Manufacturing*
- ✔ *We Stock for Same Day Shipment*
- ✔ *Five Year Warranty*
- ✔ *Free Online Demos*

Secure Console Managers

- SSHv2 Encryption
- 8, 16, or 32 DB9 Ports
- LDAP, RADIUS, TACACS+
- Internal Modem
- SNMP Monitoring
- Non-Connect Port Buffering
- SYSLOG Messages

Remote Power Switches

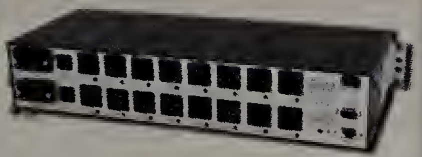

- Web Browser, Telnet and Local Access
- Dual 15 or 20 Amp Power Inputs
- Power-Up Sequencing
- Outlet Specific Passwords
- NEMA or IEC Outlets
- 4, 8, and 16 Outlet Models
- Vertical and Horizontal Models

Current Load Monitors

- Display Current/Watts/Volts
- Dual 20 Amp Circuits
- Measure Individual and Aggregate Loads
- 80% Threshold Alarm
- 120 - 208 VAC Auto Sensing
- Connect to PDUs or Reboot Switches

wti

(800) 854-7226 • www.wti.com

western telematic incorporated

5 Sterling • Irvine • California 92618-2517

Chat™|50
on the road

Chat™|150
in the office

MAXAttach™|IP
in the conference room

SIP Conference Solutions

Wherever and whenever great ideas need to be heard

ClearOne.

www.clearone.com

Introducing the Aggregator *n*TAP

www.networkTAPs.com

Efficiently aggregate full-duplex data into your analysis or security device.

- Supports 10/100/1000
- Stream into two different devices
- Rack mount up to three across
- Supports all commercial analysis systems
- Also works with open-source tools

Buffer options:

256 MB......................	**$1,495**
512 MB......................	**$1,995**

Learn more. Visit **www.networkTAPs.com.**

*n*TAP™

Choose from a variety of configurations, options, and pricing. Plus a complete line of copper and optical *n*TAPs for full-duplex analyzer systems. Free overnight delivery*

www.networkTAPs.com • 1-866-GET-*n*TAP

FC CE

*Free overnight delivery on all U.S. orders over $295 confirmed before 12 p.m. Central Time.

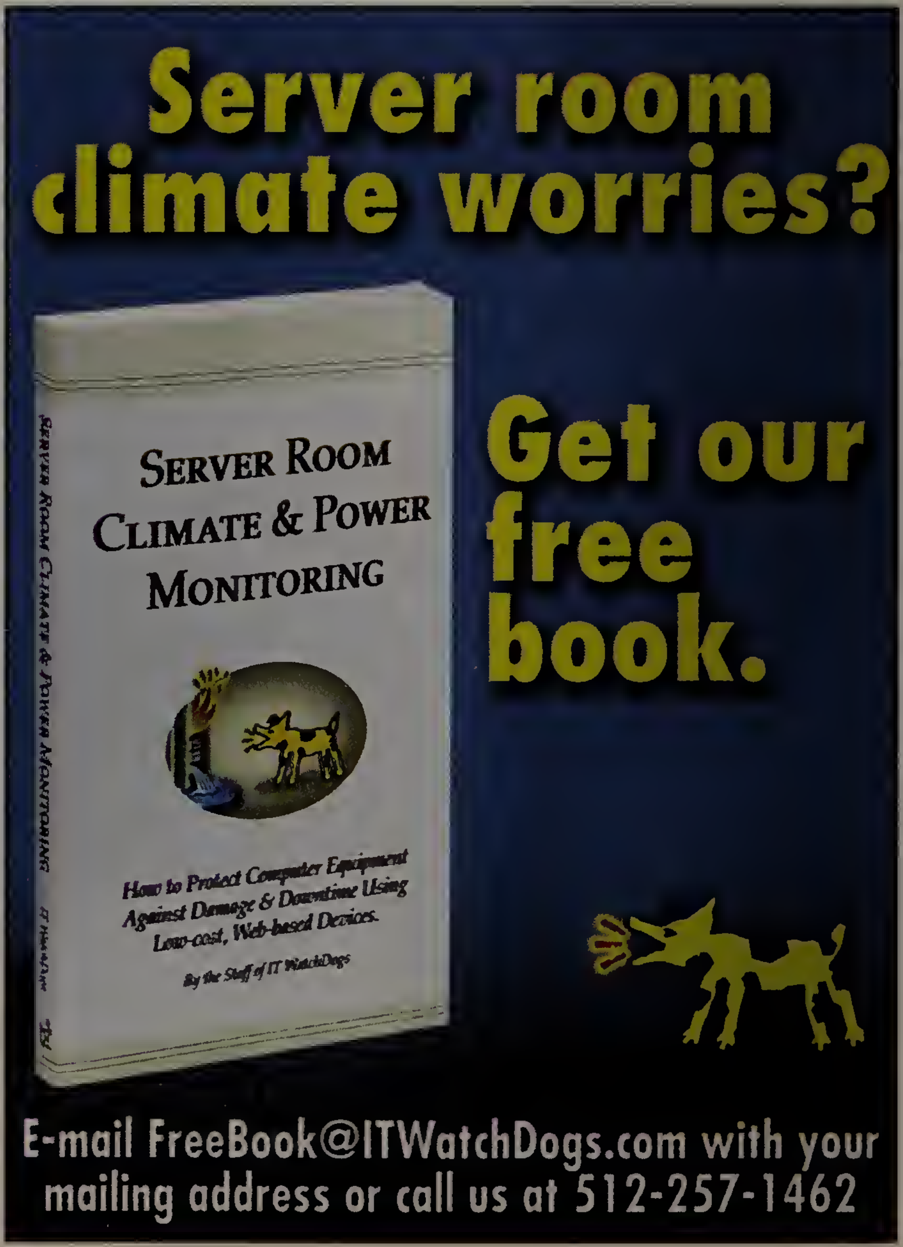

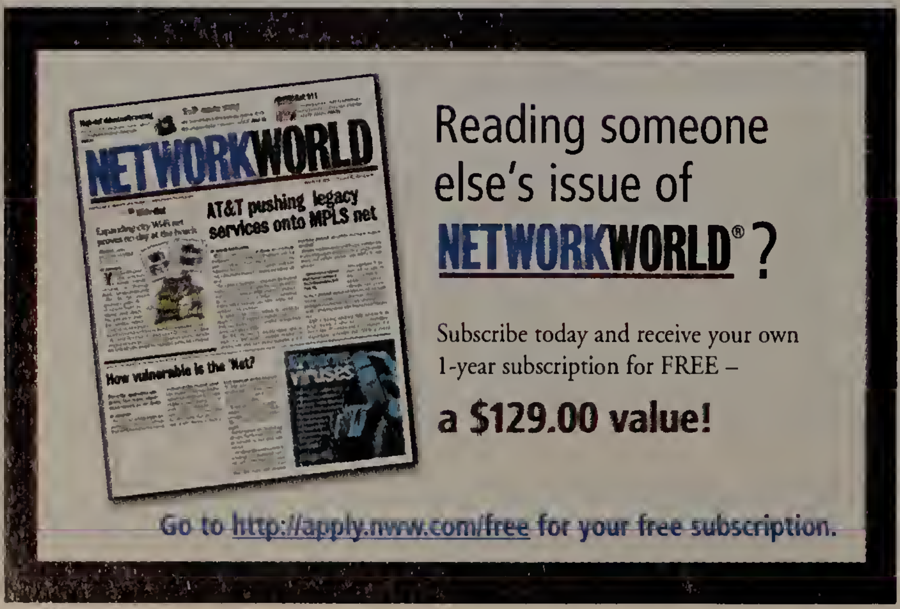

How Do You Ensure Maximum Uptime for Your Critical Devices?

With Smart Load Shedding!

Manage individual devices based on Temperature, Current Load or UPS Power Status

Should the temperature or load current exceed defined thresholds or the UPS lose power and go onto battery all or a portion of the load can be automatically shed to ensure longer operational life of your critical devices!

> **Integral Web Based GUI:** Easy-to-use, secure configuration tool

> **Remote Shutdown Agent:** Graceful server shut down and restart

> **Event Notification:** SNMP and Email alerts

> **UPS Types:** All Major UPS manufacturers supported

> **Auto-recovery:** When conditions return to normal

Server Technology
Solutions for the Data Center Equipment Cabinet

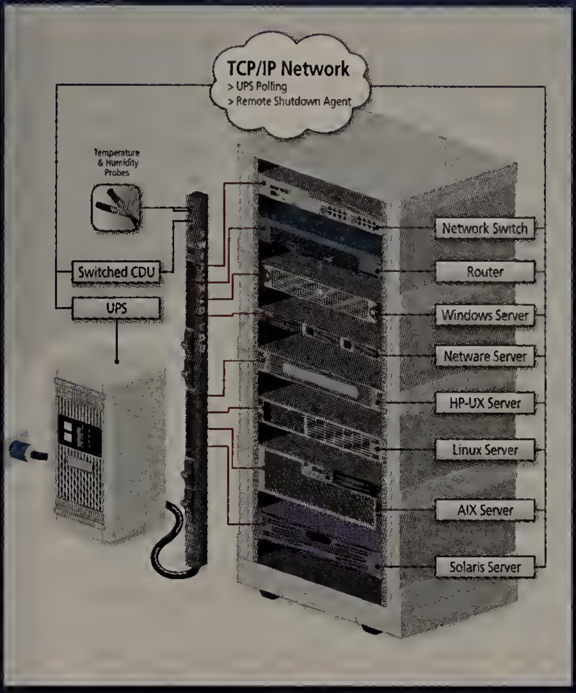

Server Technology, Inc.
1040 Sandhill Drive
Reno, NV 89521
USA

tf +1.800.835.1515
tel +1.775.284.2000
fax +1.775.284.2065

www.servertech.com
www.servertechblog.com
sales@servertech.com

SENSAPHONE
IMS-4000™

IMS SOLUTION

Monitor the REST of Your Computer Room!

- Physical Security
- Video
- Temperature
- Power Problems
- Water on the Floor
- Humidity
- Smoke and Fire
- And much more

Instant Notification by Phone or E-mail when events threaten your Infrastructure.

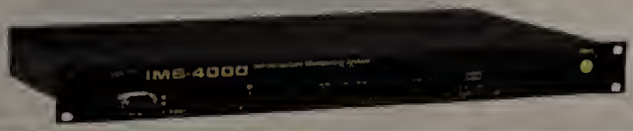

Dealers Wanted
Contact us today to discuss your application

www.ims-4000.com 877-373-2700

Let the Model 135 Monitor Your Site

The Model 135 Site Monitor is designed to serve as your "resource kit" for monitoring and maintaining computer, communications, and specialized equipment locations. With a wide range of built-in capabilities, it's easy to tailor a powerful site-specific solution.

Highlights include 10/100 Ethernet and analog modem connectivity, serial port access and text data "matching," AC and DC voltage monitoring, ping testing, and contact closure inputs and outputs. And the web-based interface makes setup and use a straight-forward process.

For complete details on the Model 135, give us a call or visit www.gkinc.com.

Gordon Kapes, Inc.
Skokie, IL USA | Ph 847-676-1750 | www.gkinc.com

HP upgrades telepresence line

BY ROBERT MULLINS

HP's improvements to its telepresence technology, unveiled last week, are aimed at making high-quality videoconferencing more accessible to business customers.

The upgrades to HP's Halo line include a less-elaborate system that produces high-definition images but at $100,000 less than the high-end model. After receiving customer feedback, HP also changed the product so a third party can join a telepresence meeting even if he is on a conventional video system or calling by phone.

The addition of new systems and features helps HP compete in a niche, but quickly growing, market for systems that make it seem as though people on each end of a videoconference are in the same room.

Telepresence system revenue, which was $64 million in 2006, will jump to $169 million in 2007 and top the $1 billion mark by 2011, according to an IDC estimate.

HP introduced Halo in December 2005. The system calls for a $349,000 telepresence studio at one location and another $349,000 studio elsewhere that looks like a mirror image of the first. A proprietary network for sending the video signal is $18,000 per month, per studio.

The new, more modest version provides high-quality video but makes it work in existing offices or conference rooms, rather than a specially made studio, for $249,000 per system.

Participants in a telepresence meeting are often taken aback by the clarity of the image from a room thousands of miles away, fed to them at 45Mbps vs. about 1Mbps on a conventional system, says Bill Wickes, director of R&D for the Halo program.

"I've seen it happen a million times by now, but people, when they first come into the room, they are startled by how realistic it is," Wickes says.

With HP's Halo telepresence system, video travels over a proprietary network at 45Mbps, vs. about 1Mbps for conventional video conferencing.

As impressive as the systems may be, demand for them is driven by more practical concerns about the time, expense, hassles and hazards of business travel.

"I call it the corporate jet replacement market," says Nora Freedman, an analyst with IDC, who adds that telepresence reduces a company's carbon footprint because forgoing air travel means fewer greenhouse gas emissions.

HP and Cisco have each introduced telepresence systems, but existing makers of conventional videoconferencing systems have a head start, Freedman says. Polycom and Teliris also offer high-definition systems that deliver a sharper image than older systems but are less expensive than the telepresence systems.

HP says it has 120 Halo studios in operation or in development worldwide.

Cisco says it has 110 of its TelePresence Systems in operation at its own facilities and has orders from 50 customers. The list price is $299,000 for a three-screen display and $79,000 for a one-screen version. ■

Cisco

continued from page 26

"That's huge for us because it allows us to focus our alarms on traffic that's specific to that segment," Nielsen says. "Currently, you have a lot of false alarms reporting" due to breaches on other segments.

For now, IronPort and its messaging security technology are the basis of SDN 3.0. There's much more to fill out, Dunlap says.

"I'd like to also know how Cisco's going to combat threats through security technology aside from e-mail security — things like IPS, behavioral technology, risk assessment," she says. "How they're going to be competing with McAfee, Symantec and others."

Phil Hochmuth, an analyst with The Yankee Group, says Cisco's emphasis on Web 2.0, collaboration and the "human network" will bring added risk as well as reward. "The next step for SDN is to deal with Web 2.0 in the enterprise, beyond the [network-centric] stuff Cisco's done well for a long time," he says.

Platon urges SDN watchers to stay tuned. "I don't think we're by any means done yet," he says. "The security issues are accelerating and the criminalization of the global network unfortunately is going to also accelerate. We're going to need to have better tools to enforce policy, better tools to understand that this is a safe connection request to accept." ■

■ **Network World**, 118 Turnpike Road, Southborough, MA 01772-9108, (508) 460-3333.

Periodicals postage paid at Southborough, Mass., and additional mailing offices. Posted under Canadian International Publication agreement #41321518. Network World (ISSN 0887-7661) is published weekly, except for combined issues for the first two weeks of July and the last two weeks in December by Network World, Inc., 118 Turnpike Road, Southborough, MA 01772-9108.

Network World is distributed free of charge in the U.S. to qualified management or professionals.

To apply for a free subscription, go to www.subscribenw.com or write Network World at the address below. No subscriptions accepted without complete identification of subscriber's name, job function, company or organization. Based on the information supplied, the publisher reserves the right to reject non-qualified requests. Subscriptions: 1-508-490-6444.

Nonqualified subscribers: $5.00 a copy; U.S. - $129 a year; Canada - $160.50 (including 7% GST, GST#126659952); Central & South America - $150 a year (surface mail); all other countries - $300 a year (airmail service). Four weeks notice is required for change of address. Allow six weeks for new subscription service to begin. Please include mailing label from front cover of the publication.

Network World can be purchased on 35mm microfilm through University Microfilm Int., Periodical Entry Dept., 300 Zebb Road, Ann Arbor, Mich. 48106.

POSTMASTER: Send Change of Address to **Network World**, P.O. Box 3090, Northbrook, IL 60065. Canadian Postmaster: Please return undeliverable copy to PO Box 1632, Windsor, Ontario N9A7C9.

Reprints (minimum 500 copies) and permission to reprint may be purchased from Reprint Management Services at (717) 399-1900 x128 or networkworld@reprintbuyer.com.

USPS735-730

NETWORKWORLD

Editorial Index

Advertiser Index

Advertiser	Page #	URL
American Power Conversion	23, 48	www.apc.com
Berkeley Varitronics Systems	50	www.bvsystems.com
CA Inc	9	ca.com/secure
ClearOne Communications Inc	49	www.clearone.com
Diskeeper Corp	55	www.diskeeper.com/nw7
DNSstuff LLC	25	DNSstuff.com
dtSearch Corp	40	www.dtsearch.com
Gordon Kapes Inc	51	www.gkinc.com
Hewlett Packard	11	www.procurve.com/proactive
Hewlett Packard	56	www.hp.com/go/breakthecycle71
IBM Corp	28-29	ibm.com/takebackcontrol/ready
IBM Corp	32-33	ibm.com/takebackcontrol/connect
IBM Corp	45, 47	ibm.com/takebackcontrol/green
IT Watchdogs	50	ITWatchdogs.com
Intel Corp	5, 7	Intel.com/xeon
Juniper Networks Inc	13	www.juniper.net/threatmanagement
Microsoft Corp	2-3	microsoft.com/voip
NEC Corp	17	www.necus.com/necip
NEC Corp	19	www.necus.com/digitalsignage
NEC Corp	21	www.necus.com/vpcc
Network Instruments LLC	27	www.NetworkInstruments.com/TimeTravel
Network Instruments LLC	49	www.networkTAPs.com
Quadrics Inc	15	www.quadrics.com
Rose Electronics	50	www.rose.com
Sensaphone	51	www.ims-4000.com
Server Technology Inc	51	www.servertech.com
Western Telematic Inc	49	www.wti.com

These indexes are provided as a reader service. Although every effort has been made to make them as complete as possible, the publisher does not assume liability for errors or omissions.

*Indicates Regional Demographic

IDG

Patrick J. McGovern, Chairman of the Board
Bob Carrigan, President, IDG Communications

Network World is a publication of IDG, the world's largest publisher of computer-related information and the leading global provider of information services on information technology. IDG publishes over 300 computer publications in 85 countries. One hundred million people read one or more IDG publications each month. *Network World* contributes to the IDG News Service, offering the latest on domestic and international computer news.

Network World, Inc.

118 Turnpike Road, Southborough, MA 01772
Phone: (508) 460-3333

TO SEND E-MAIL TO NWW STAFF
firstname_lastname@nww.com

CEO/Publisher: Evilee Ebb
President/General Manager, Online: John Gallant
Chief Operating Officer: W. Michael Draper

ONLINE SERVICES
V. P., Online Sales/Operations: Susan Cardoza
V. P., Online Audience/Architecture: Dan Gallagher
V. P., Online Development/Custom Media: Dan Hirsh

SEMINARS, EVENTS AND IDG EXECUTIVE FORUMS
Exec. V. P., Events/Executive Forums: Neal Silverman
Vice President, Event Sales: Andrea D'Amato
Sr. Director, Marketing/Business Dev.: Mike Garity

MARKETING
Director of Marketing: Donna Pomponi

AD OPERATIONS
Production Manager: Jami Thompson
Advertising Coordinator: Maro Eremyan

FINANCE
Vice President Finance: Mary Fanning

HUMAN RESOURCES
Human Resources Manager: Eric Cormier

CIRCULATION
Senior Director of Circulation: Richard Priante

INFORMATION SERVICES
Chief Operating Officer: Mike Draper
Director of Systems Development: Tom Kroon
Director of Network IT: Rocco Bortone
Manager of Database Technologies: Puneet Narang

IDG LIST RENTAL SERVICES
Director of List Management, Steve Tozeski
Toll free: (800) IDG-LIST (US only)/Direct: (508) 370-0822

Sales Offices

New York/New Jersey
Tom Davis, Associate Publisher, Eastern Region
Elisa Della Rocco, Regional Account Director
Internet: tdavis, elisas@nww.com
(201) 634-2300/FAX: (201) 634-9286

Northeast
Elisa Della Rocco, Regional Account Director
Internet: elisas@nww.com
(201) 634-2300/FAX: (201) 634-9286

Mid-Atlantic
Jacqui DiBianca, Regional Account Director
Internet: jdibian@nww.com
(610) 971-0808/FAX: (201) 621-5095

Midwest/Central
Tom Davis, Associate Publisher, Eastern Region
Internet: tdavis@nww.com
(201) 634-2300/FAX: (201) 634-9286

Southeast
Don Seay, Regional Account Director
Internet: dseay@nww.com
(404) 874-7232/FAX: (201) 621-6129

Northern California/Northwest
Sandra Kupiec, Assoc. Publisher, Western Region
Karen Wilde, Regional Account Director
Vanessa Tormey, Regional Account Director
Coretta Wright, Regional Account Manager
Katie Layng, Account Executive
Internet: skupiec, kwilde, vtormey, cwright, klayng, jhallett, ckim, tmarfori@nww.com
(415) 243-4122/FAX: (415) 267-4519

Southwest/Rockies
Becky Bogart, Regional Account Director
Internet: bbogart@nww.com
(949) 713-5153/FAX: (949) 334-1155

Online/Integrated Solutions
Susan Cardoza, V. P. Online Sales/Operations
Debbie Lovell, Sr. Online Account Manager
Denise Landry, Sr. Account Coordinator
Justine Wallace, Account Coordinator
Internet: scardoza, dlovell, dlandry, jwallace@nww.com
(508) 460-3333/FAX: (508) 490-6603

Marketplace / Emerging Markets
Enku Gubaie, Mgr. of Marketplace/Emerging Markets
Internet: egubaie@nww.com
(508) 460-3333/FAX: (508) 460-1192

BACKSPIN
Mark Gibbs

Waxing philosophical about failure modes

A couple of weeks ago Mich Kabay wrote about an article in the *Wall Street Journal* that discussed, albeit in a lame, noob kind of way, techniques for employees "to get around the IT departments."

Through a curious process understandable only by those with a Ph.D. in quantum mechanics and its relationship to the publishing process Kabay's newsletter article was attributed to me.

As much as I enjoyed this I had to confess to all who wrote in to compliment me on my (Kabay's) article that I was an innocent bystander splattered with the mud of someone else's writing. I wish such a mistake occurred more often with the checks coming to my address.

Be that as it may, a letter of a complimentary nature raised an interesting question. One of my (Kabay's) readers wrote in to ask: "Have you considered the other perspective on the WSJ article, namely the full disclosure/'king hath no clothes' side to it? If desktop systems were actually secured (or for that matter fully securable), the holes would not be there to be exploited by sneaky employees, nor to be exposed by WSJ."

The second issue of realistically and comprehensively securing desktop systems is, regrettably, implausible. If you take any reasonably complex system (no, I am not about to define "reasonably", just work with me people) then it is obvious that when flying pigs deliver such systems life will be vastly improved. Until that day we have to live with systems that are secured subject to two limitations: What we know and what we can afford.

What we know about any complex system is always limited because of Turing's Halting Problem. In a roundabout sort of way, this says that figuring out by inspection of computer code whether a particular state, such as stopping or deducing the existence of rice pudding from first principles, can occur is impossible.

This means that where we're considering security, then identifying and characterizing all failure modes (aka security problems) also is impossible. Even worse, identifying most failure modes is equally impossible because we can't know the limits of what we don't know because Turing says so.

What we can afford over the short term is the discovery of the most easily found failure modes. These modes are those that are easily and therefore inexpensively found (as in a few dollars each). To identify the next set of failure modes that are harder to find is more expensive and so on until we are spending the equivalent of the gross domestic product of Bolivia to find a single failure mode.

Even worse is the problem that there is no correlation between how many failure modes we know of and what it might cost to find half of them. We can't determine what kind of cost will be involved we try to remove as many vulnerabilities as possible from a system. What we can be sure of is that it can't be done completely.

But it is our reader's first question on the value of full disclosure that is the most interesting. It is obvious that the most common state of knowledge about failure modes is when both the good guys and the bad guys don't know a failure mode exists. This is good because the good guys are at no disadvantage. But what of the other situations when either or both parties know that a failure mode exists? We'll return to that next week.

Gibbs waxes philosophical from Ventura, Calif. Send your theories to backspin@gibbs.com.

Michael Cooney

LAYER 8
A coffee break for your head

Is the bloom off municipal Wi-Fi?

With last week's news that Chicago and San Francisco are blowing up their wireless broadband plans, the bloom is blown off the municipal Wi-Fi movement.

Chicago cited rising costs, spotty demand and uncooperative carriers as the main reasons for the cancellation of the $18.5 million rollout that would have covered the city's 228 square miles. According to one report, EarthLink and AT&T demanded that Chicago become an "anchor tenant," paying an annual fee to use the Wi-Fi network to support city services. When the city refused — and insisted that the system attached to city street lights and lamp poles be built, maintained and operated at the contractor's "sole expense" — the whole system came crashing down.

Meanwhile EarthLink's contract to build a municipal Wi-Fi network in San Francisco appears to be dead following a restructuring of the struggling ISP. Just last week, EarthLink said it would cut 900 jobs and shutter several regional offices. "We will not devote any new capital to the old muni Wi-Fi model that has us taking all of the risk by fronting all of the capital, then paying to buy our customers one by one," said EarthLink President and CEO Rolla Huff. That includes currently planned networks where the company hasn't yet made capital investments, meaning San Francisco and other cities, would have to build municipal Wi-FI networks under other arrangements.

EarthLink has won contracts for networks in Houston; Corpus Christi, Texas; and other cities. *The Houston Chronicle* reported that EarthLink is months behind schedule in getting started with Houston's Wi-Fi project, and there are doubts it will go forward at all.

Published reports say Chicago and San Francisco are the latest in a string of municipalities to encounter troubles with their municipal broadband initiatives because of ballooning budgets and dwindling usage. Anchorage, Alaska, and Corona, Calif., have discontinued their municipal wireless projects after MetroFi, the private industry partner in both cities, said it could not offer free service without a commitment from each municipality to be an anchor tenant. About 175 U.S. cities or regions have citywide or partial systems.

However, Philadelphia, for example, is well on its way to becoming one of the world's biggest Wi-Fi hot spots. *Network World* recently reported that Philadelphia gave EarthLink the green light to cover the 135-square-mile city with a wireless mesh network by year-end. EarthLink is moving full-speed ahead, adding Tropos Networks access points to light poles around the city, testing and optimizing the network, and building out coverage at a pace of 5,000 potential households per workday. Today, coverage has expanded to 80% of the city.

Pssst . . . Wanna buy a data center?

So what do you do with 250 servers and thousands of terabytes of data storage when nobody else wants it? Auction it online — what else? High-tech online asset liquidator Rasmus Auctioneers is prepping $15 million worth of brand-new — still in the box — data center gear that was dumped in its lap from a Department of the Interior lease cancellation. The entire lot, which includes Egenera blade servers, EMC Centera Servers and ADIC Digital Tape Libraries, is online to be sold to the highest bidder.. The inventory will be sold by Internet-only auction at 2 p.m. EST on Sept. 12. "The liquidation will be like an eBay sale on steroids," Rasmus said in a statement. The possible fly-in-the-ointment with this data center gear is that none of it comes with licenses, support or extended warranties, Rasmus said. Still, the equipment could fill a variety of roles from helping a company branch out their current data center or improving redundancy and backup systems.

Cooney is filling in for Paul McNamara, who is on vacation. He can be reached at mcooney@nww.com.

Advertisement

The Key to Maximum Performance and Reliability for Windows Vista™ and Beyond

A SPECIAL REPORT

Windows Vista has finally arrived, and reviewers are hailing it as the best OS Microsoft® has ever built. For corporations, it boasts robust features such as greatly improved security, a wholly new and highly versatile user interface, significantly simplified software deployment, and broad backwards-compatibility. While it may not happen right away, most if not all Windows®-central enterprises will want to make the move to Windows Vista.

Defragmentation Technology—Time for a Change

It's well known that a high number of system slows, crashes, and even file corruption and errors can be traced to file fragmentation.[1] File fragmentation puts your system performance and reliability in serious jeopardy. It's no surprise, then, that substantial performance gains from defragmenting, in the range of 90%, have been documented.[2]

But it's not only the decision to defragment your systems that makes the difference. The choice of defragmentation technology, both before and after your move to Windows Vista, is crucial.

The sheer scope and activity of computer systems today has made even scheduled defragmentation, once "state of the art," obsolete. Disks and files once measured in kilobytes and megabytes are now measured in gigabytes and terabytes, and the sheer number of files has increased tremendously. Testing has shown that scheduled defragmentation cannot keep pace; between defragmenter runs, fragmentation simply builds up and continues to negatively impact performance.[3]

The True Solution to Maximum Performance and Reliability

Only a completely automatic defragmentation solution such as Diskeeper 2007—released just in time for Windows Vista—will truly keep pace with the ever-expanding capacity and intense activity on today's disks. Instead of providing partial benefit when defragmentation runs occur, *all* applications and *all* files benefit from increased performance *all* the time.

With its proprietary breakthrough InvisiTasking™ technology, Diskeeper 2007 defragments and enhances file systems in real-time, with no scheduling needed. Defragmentation is now performed on-the-fly, with no performance hit on system resources. Your system is consistently faster and more reliable with Diskeeper 2007—period. In testing against scheduled defragmentation, which leaves fragmented files behind after running, Diskeeper 2007 consistently *eliminates* fragmentation to continuously provide maximum performance and reliability.[3] Take advantage of our free 45-day trial and see for yourself.

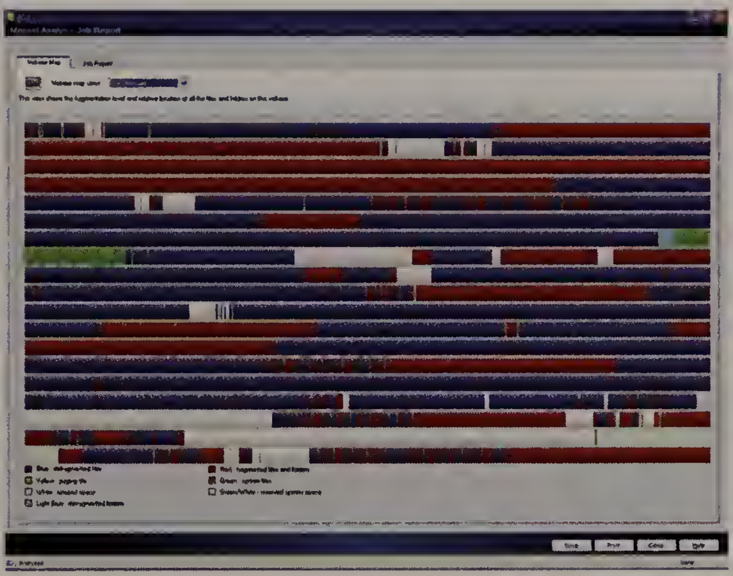

Diskeeper's interface shows fragmentation levels and relative locations of all the files and folders on the selected volume.

Plus, Diskeeper 2007 includes Intelligent File Access Acceleration Sequencing Technology (I-FAAST™) 2.0, specifically designed to deliver increased performance, speed and reliability above and beyond defragmentation benefits.

Be Completely Ready for Windows Vista

With its stunning GUIs, Windows Vista brings a whole new level of operation to computer interaction. Because of its graphical nature, and its support of an ever-widening variety of graphical and video-based programs, enormous files and high-capacity disks are the norm. Smooth, fast access to these files is vital, especially with applications such as business conferencing and video presentations. Additionally, Windows Vista utilizes considerable resources, and it is vital that applications offering better performance not drain resources from an already taxed pool.

If scheduled defragmentation cannot keep up with current system demands, it will be completely lost with Windows Vista. Deploying Diskeeper's real-time defragmentation right at Vista deployment means that peak performance and reliability are part of the package, and one less headache for an already-overworked system administrator.

With Windows Vista, disk activity on servers also reaches new demanding heights—and Diskeeper Server and Diskeeper EnterpriseServer versions are right there with advanced technologies such as Terabyte Volume Engine™ 2.0, especially designed for fast defragmentation on the highest capacity servers.

Diskeeper's automatic defragmentation is vital *during* the move to Windows Vista as well. Deployment of a new OS is no mean feat—it means hardware upgrades, changes and revisions in policy, verification of legacy support, carefully controlled software deployment, and a long list of other vital tasks performed while continually extinguishing fires and maintaining current networks. The last things you need during such an evolution are reliability and performance problems from your current systems such as slowed disk access and response times.

And since Diskeeper 2007 already runs on Windows Vista, the licenses you buy now will be with you every step of the way, all the way into and beyond the move to Windows Vista.

Automatically maximize your system performance and reliability *now* and put disk performance problems behind you—for Windows Vista and beyond.

Diskeeper® 2007
Maximizing Performance and Reliability
—*Automatically!*™

Special Offer

Try Diskeeper 2007 FREE for 45 days!
Download: www.diskeeper.com/nw7
(Note: Special 45-day trialware is only available at the above link)

Volume licensing and Government / Education discounts are available from your favorite reseller or call 800-829-6468 code 4413

For test results, white papers and case studies, visit www.diskeeper.com/NWpaper

[1] File Fragmentation White Paper
[2] Article: *The Impact of Disk Fragmentation, WindowsITPro*
[3] White Paper: *Is Real-Time Defrag Needed?*

Diskeeper corporation

SERVER ROOM
SERVER ROOM
SERVER ROOM
SERVER ROOM